arXiv:1611.06501v1 [cs.DS] 20 Nov 2016

Approximation and parameterized algorithms for geometric independent set with shrinking*

Michał Pilipczuk† Erik Jan van Leeuwen‡ Andreas Wiese§

Abstract

Consider the Maximum Weight Independent Set problem for rectangles: given a family of weighted axis-parallel rectangles in the plane, find a maximum-weight subset of non-overlapping rectangles. The problem is notoriously hard both in the approximation and in the parameterized setting. The best known polynomial-time approximation algorithms achieve super-constant approximation ratios [6, 9], even though there is a $(1+\epsilon)$-approximation running in quasi-polynomial time [2, 10]. When parameterized by the target size of the solution, the problem is W[1]-hard even in the unweighted setting [22].

To achieve tractability, we study the following *shrinking model*: one is allowed to shrink each input rectangle by a multiplicative factor $1-\delta$ for some fixed $\delta > 0$, but the performance is still compared against the optimal solution for the original, non-shrunk instance. We prove that in this regime, the problem admits an EPTAS with running time $f(\epsilon,\delta)\cdot n^{\mathcal{O}(1)}$, and an FPT algorithm with running time $f(k,\delta)\cdot n^{\mathcal{O}(1)}$, in the setting where a maximum-weight solution of size at most k is to be computed. This improves and significantly simplifies a PTAS given earlier for this problem [1], and provides the first parameterized results for the shrinking model. Furthermore, we explore kernelization in the shrinking model, by giving efficient kernelization procedures for several variants of the problem when the input rectangles are squares.

1 Introduction

The classic Maximum (Weight) Independent Set problem is defined as follows: given a graph, the goal is to select a set of pairwise non-adjacent vertices with maximum cardinality or maximum total weight. In its full generality, the problem is NP-hard and intractable from the perspective of both approximation and parameterized algorithms: it is NP-hard to approximate within ratio $n^{1-\epsilon}$ for any $\epsilon > 0$ [32], and it is W[1]-hard when parameterized by the solution size [13]. Therefore, many restricted settings have been studied.

One well-studied case is to consider a geometric setting. That is, the input consists of a family of geometric objects, and the goal is to select a maximum-weight subfamily of pairwise non-overlapping objects. This case reduces to the graph setting by considering the intersection graph of the objects. The graphs that arise in this way are highly structured, which gives hope for better results than for general graphs.

This paper concentrates on the variant in which the input objects are axis-parallel rectangles in the two-dimensional plane. In this variant, Maximum Weight Independent Set admits much

*The research of Mi. Pilipczuk is supported by Polish National Science Centre grant UMO-2013/11/D/ST6/03073. Mi. Pilipczuk is also supported by the Foundation for Polish Science (FNP) via the START stipend programme.

†Institute of Informatics, University of Warsaw, Poland, michal.pilipczuk@mimuw.edu.pl.

‡MPI für Informatik, Saarland Informatics Campus, Saarbrücken, Germany, erikjan@mpi-inf.mpg.de.

§Departamento de Ingenieria Industrial, Universidad de Chile, Chile, awiese@dii.uchile.cl.

smaller approximation ratios than on general graphs. There is an $(1+\epsilon)$-approximation algorithm running in polynomial time (PTAS) if all input objects are squares [7, 14, 28], an $\mathcal{O}(\log\log n)$-approximation algorithm for unweighted rectangles [6], and an $\mathcal{O}(\log n/\log\log n)$-approximation algorithm for weighted rectangles [9]. Recently, there has been much interest in designing *quasi-polynomial approximation schemes* (*QPTASes*), that is, $(1+\epsilon)$-approximation algorithms running in quasi-polynomial time. This research direction culminated in giving a QPTAS for the case of (not necessarily convex) polygons in the plane [3, 17]. Whether this can be improved to a PTAS remains a challenging open problem.

From the parameterized perspective, the problem remains W[1]-hard when parameterized by the k size of the solution, even for unweighted unit squares [22]. Therefore, the existence of an FPT algorithm with running time $f(k)\cdot n^{\mathcal{O}(1)}$ for a computable f is unlikely under standard assumptions from parameterized complexity; this also excludes the existence of an EPTAS for the problem [22]. However, the problem admits a faster-than-brute-force parameterized algorithm with running time $n^{\mathcal{O}(\sqrt{k})}$, which is optimal under the Exponential Time Hypothesis [23]. This algorithm works in the general setting of finding a maximum-weight independent set of size k in a family of polygons in the plane.

Shrinking model. In order to circumvent some of the many challenges that arise when designing approximation or parameterized algorithms for geometric Maximum Weight Independent Set, we investigate the *shrinking model* introduced by Adamaszek et al. [1]. In this model, one is allowed to shrink each input object by a multiplicative factor $1-\delta$ for some fixed $\delta > 0$, but the weight of the computed solution is still compared to the optimum for the original, non-shrunk instance; see Section 2 for a formal definition. It is known that the shrinking model allows for substantially better approximation algorithms than the general setting: Adamaszek et al. [1] gave a PTAS for axis-parallel rectangles, which was later generalized by Wiese to arbitrary convex polygons [30]. However, it has not been studied so far whether shrinking also helps to design parameterized algorithms. One concrete question would be whether Independent Set for axis-parallel rectangles remains W[1]-hard in the shrinking model.

Our results. This paper addresses the parameterized complexity of Maximum Independent Set of Rectangles in the shrinking model, and answers the above questions positively. On the way to our two main parameterized contributions, we also improve the PTAS by Adamaszek et al. [1] to an EPTAS.

Our first main contribution is that Maximum Independent Set of Rectangles is fixed-parameter tractable (FPT) in the shrinking model. More precisely, for a shrinking parameter δ, we can decide in randomized time $f(k,\delta)\cdot(nN)^{\mathcal{O}(1)}$ whether there is an independent set of k (shrunk) rectangles, or the original family has no independent subfamily of size k. Here, N is the total bit size of the input and f is some computable function. The algorithm works also in the weighted setting, where we are looking for a maximum-weight subset of at most k non-overlapping rectangles. The main reason why we are able to circumvent the W[1]-hardness for the standard model (i.e., without shrinking) is that the reduction of Marx [22] heavily relies on tiny differences in the coordinates of the rectangles. However, as Adamaszek et al. [1] and this paper show, this aspect vanishes in the shrinking model.

The parameterized algorithm is actually a consequence of an EPTAS that we present for Maximum Weight Independent Set of Rectangles. That is, we give an algorithm with running time $f(\epsilon,\delta)\cdot(nN)^{\mathcal{O}(1)}$ that finds a subset of rectangles that do not overlap after shrinking by factor $1-\delta$, and whose total weight is at least $1-\epsilon$ times the optimum without shrinking. Recall that the standard model does not admit an EPTAS, unless $\mathsf{FPT} = \mathsf{W}[1]$ [22].

Our EPTAS is based on the same principles as the PTAS of Adamaszek et al. [1]. The idea is

to assemble an optimum solution using a bottom-up dynamic-programming approach pioneered by Erlebach et al. [14]. Each subproblem solved in the dynamic program corresponds to the maximum weight of an independent set contained in a "box", and the computation of the optimum for each such box boils down to enumerating a limited number of carefully chosen partitions of the box into smaller boxes. Intuitively, the ability to shrink is used to make sure that rectangles fit nicely into the different boxes. The main challenge is to ensure that these boxes can be assumed to be simple, and therefore only a limited number of subproblems is necessary to assemble a near-optimum solution.

The crucial contribution in our approximation algorithm is that we show that rectangular boxes suffice. In [1], most rectangles were shrunk in *only one direction* and therefore, the boxes were axis-parallel polygons with at most $g(\epsilon, \delta)$ sides each, for some function g. This makes the dynamic program very complex, and yields a running time of $(nN)^{g(\epsilon,\delta)}$ due to the sheer number of subproblems solved. In this paper, we fully exploit the properties of the shrinking model and shrink *each* rectangle in *two directions*. This changes the analysis, but the main advantage is that we only need to consider boxes that are rectangles (i.e., with only four sides) in our dynamic program. This greatly simplifies the dynamic program, and we show that we need to consider only $f(\epsilon, \delta) \cdot (nN)^{\mathcal{O}(1)}$ different subproblems. Hence, our EPTAS is both substantially faster and significantly simpler than previous work.

Our second main contribution shows that several important subcases of Maximum Weight Independent Set of Rectangles with δ-shrinking admit a polynomial kernel when parameterized by k and δ. Intuitively, such a kernel is a polynomial-time algorithm that finds a subfamily of the input rectangles of size bounded by a polynomial of k and δ that retains an optimum solution after δ-shrinking; a formal definition is given in Section 4. We show that for unit squares of non-uniform weight, we can construct a kernel of size $\mathcal{O}(k/\delta^2)$, while for squares of non-uniform size, but of uniform weight, we can construct a kernel of size $\mathcal{O}(k^2 \cdot \frac{\log(1/\delta)}{\delta^3})$. As a direct consequence, we obtain FPT algorithms for the considered variants with running time $(k/\delta)^{\mathcal{O}(\sqrt{k})} \cdot (nN)^{\mathcal{O}(1)}$ by applying the $n^{\mathcal{O}(\sqrt{k})}$-time algorithm of Marx and Pilipczuk [23] on the kernels. This *subexponential* running time is far better than the running time of the FPT algorithm for the general rectangle case. Other results from the literature, like Alber and Fiala [4], Hunt et al. [19], or Chan et al. [8], can be also used in combination with our kernelization to yield faster FPT and approximation algorithms.

The main idea of our kernelization algorithms is as follows. If two or more squares have similar size and their centers lie very close to each other, then in the shrinking model we can replace them by a single (shrunk) square that is contained in each of them. In the weighted setting, the new square keeps the largest among the weights of the original ones. Thus, we can assume that squares of roughly the same size do not lie too close to each other. This enables a greedy approach, similar in spirit to Alber and Fiala [4], which works for unit squares both in the uniform and non-uniform weight setting. For squares of non-uniform size, but uniform weight, we observe that if two squares have very different size and they overlap, then either one is contained in the other (and thus we can remove the larger one without changing the optimum), or after the shrinking they do not overlap anymore. Hence, if there are k squares of pairwise very different size, then they must be non-overlapping after shrinking, and we have a solution. Otherwise, the squares "occupy" only $\mathcal{O}_\delta(k)$ levels of magnitude, and we can kernelize each of these levels separately.

Organization. In Section 2 we establish the notation. Section 3 contains the presentation of these results, with the FPT algorithm presented as an adjustment of the EPTAS. In Section 4 we explore kernelization for the considered problems; this section also contains the formal statements of the kernelization results. Finally, in Section 5 we gather some concluding remarks and provide an outlook on future work.

2 Preliminaries

For clarity, we essentially adopt the notation of Adamaszek et al. [1]. Suppose that $\mathcal{R} = \{R_1, R_2, \ldots, R_n\}$ is a family of axis-parallel rectangles given in the input. Each rectangle R_i is described as

$$R_i = \{(a,b) : x_i^{(1)} < a < x_i^{(2)} \text{ and } y_i^{(1)} < b < y_i^{(2)}\},$$

where $x_i^{(1)} < x_i^{(2)}$ and $y_i^{(1)} < y_i^{(2)}$ are integers. Thus, the input rectangles are assumed to be open, and their vertices are at integral points. In general, this is not a restriction [21, Lemma 2.1]. We assume that the family $\mathcal{R}$ is given in the input with all the coordinates $x_i^{(1)}, x_i^{(2)}, y_i^{(1)}, y_i^{(2)}$ encoded in binary; thus, the coordinates are at most exponential in the total bit size of the input, denoted by N. For a rectangle R_i, we define its *width* $g_i = x_i^{(2)} - x_i^{(1)}$ and *height* $h_i = y_i^{(2)} - y_i^{(1)}$. Moreover, each rectangle R_i has a prescribed *weight* w_i, which is a nonnegative real. For a subset $\mathcal{S} \subseteq \mathcal{R}$, we denote $w(\mathcal{S}) = \sum_{R_i \in \mathcal{S}} w_i$.

Fix a constant δ with $0 < \delta < 1$. For a rectangle R_i, its *δ-shrinking* $R_i^{-\delta}$ is the rectangle with x-coordinates $x_i^{(1)} + \frac{\delta}{2} g_i$ and $x_i^{(2)} - \frac{\delta}{2} g_i$, and y-coordinates $y_i^{(1)} + \frac{\delta}{2} h_i$ and $y_i^{(2)} - \frac{\delta}{2} h_i$. The δ-shrinking retains the weight w_i of the original rectangle. For a subset $\mathcal{S} \subseteq \mathcal{R}$, we denote $\mathcal{S}^{-\delta} = \{R_i^{-\delta} : R_i \in \mathcal{S}\}$ to be the family of δ-shrinkings of rectangles from $\mathcal{S}$.

A family of rectangles is *independent* (or is an *independent set*) if the rectangles are pairwise non-overlapping. In the Maximum Weight Independent Set of Rectangles problem (MWISR) we are given a family of axis-parallel rectangles $\mathcal{R} = \{R_1, R_2, \ldots, R_n\}$, and the goal is to find a subfamily of $\mathcal{R}$ that is independent and has maximum total weight. This maximum weight will be denoted by $\mathsf{OPT}(\mathcal{R})$. In the parameterized setting, the *parameterized* MWISR problem, we are additionally given an integer parameter k, and we look for a subfamily of $\mathcal{R}$ that has size at most k, is independent, and has maximum possible weight subject to these conditions. This maximum weight will be denoted by $\mathsf{OPT}_k(\mathcal{R})$.

In the δ-shrinking setting, we relax the requirement of independence to just requiring the disjointness of δ-shrinkings, but we still compare the weight of the output of our algorithm with $\mathsf{OPT}(\mathcal{R})$, respectively with $\mathsf{OPT}_k(\mathcal{R})$.

Finally, given a family $\mathcal{R}$ of arbitrary objects in the plane, the *intersection graph* G induced by $\mathcal{R}$ has a vertex for each object in $\mathcal{R}$ and an edge between two vertices if and only if the corresponding objects intersect. Conversely, the family $\mathcal{R}$ is said to be a *representation* of G.

3 Main results

With the above definitions in mind, we can state formally our main results.

Theorem 1 (FPT for MWISR with δ-shrinking)**.** *There is a randomized algorithm that given a weighted family $\mathcal{R}$ of n axis-parallel rectangles with total encoding size N, and parameters k and δ, runs in time $f(k,\delta) \cdot (nN)^c$ for some computable function f and constant c, and outputs a subfamily $\mathcal{S} \subseteq \mathcal{R}$ such that $|\mathcal{S}| \leq k$, $\mathcal{S}^{-\delta}$ is independent, and $w(\mathcal{S}) \geq \mathsf{OPT}_k(\mathcal{R})$ with probability at least $1/2$.*

Theorem 2 (EPTAS for MWISR with δ-shrinking)**.** *There is a randomized algorithm that given a weighted family $\mathcal{R}$ of n axis-parallel rectangles with total encoding size N, and parameters δ, ϵ, runs in time $f(\epsilon,\delta) \cdot (nN)^c$ for some computable function f and constant c, and outputs a subfamily $\mathcal{S} \subseteq \mathcal{R}$ such that $\mathcal{S}^{-\delta}$ is independent, and $w(\mathcal{S}) \geq (1-\epsilon)\mathsf{OPT}(\mathcal{R})$ with probability at least $1/2$.*

We prove the above thereoms in this section. We concentrate on the proof of Theorem 2, because Theorem 1 follows by a simple adjustment of our argumentation, as we will explain in Section 3.3.

Throughout the proof we fix the input family $\mathcal{R} = \{R_1, \ldots, R_n\}$, and we denote $\mathsf{OPT} = \mathsf{OPT}(\mathcal{R})$. We also fix the constants δ and ϵ, and w.l.o.g. we assume that $1/\delta$ and $1/\epsilon$ are even integers larger than 4. For convenience, throughout the proof we aim at finding a solution $\mathcal{S}$ with $w(\mathcal{S}) \geq (1 - d \cdot \epsilon)\mathsf{OPT}$ for some constant d, for at the end we may rescale the parameter ϵ to ϵ/d.

Recall that N denotes the total size of the binary encoding of all the coordinates of the input rectangles. By shifting all the rectangles, we may assume without loss of generality that they all fit into the square $[1, L] \times [1, L]$, where $L = (1/\delta\epsilon)^\ell$ for some integer $\ell = \mathcal{O}(N)$. That is, all the coordinates $x_i^{(1)}, x_i^{(2)}, y_i^{(1)}, y_i^{(2)}$ are between 1 and L, so in particular the width and the height of each rectangle is smaller than L.

Our reasoning is divided into two main steps. First, like in the PTAS due to Adamaszek et al. [1], in Section 3.1 we use standard shifting arguments to remove some rectangles from $\mathcal{R}$ so that OPT decreases only by an $\mathcal{O}(\epsilon)$-fraction, but the resulting family admits some useful properties. Then, we shrink the rectangles in a similar way as in [1]. As we will point out, there is a subtle but important difference in our shrinking procedure compared to [1]. Second, we show that the properties of the obtained family enable us to compute an optimum solution using dynamic programming; this algorithm is presented in Section 3.2.

3.1 Sparsifying the family

Separating rectangles by size. Intuitively, we will apply shifting techniques to extract some structure in the input family $\mathcal{R}$ while losing only an $\mathcal{O}(\epsilon)$-fraction of OPT. The first goal is to obtain a classification of the rectangles according to their widths and heights, such that rectangles in the same class have similar widths and heights, but between the classes the dimensions differ significantly. This is encapsulated formally in the following definition.

Definition 1. A subfamily $\mathcal{R}' \subseteq \mathcal{R}$ is *well-separated* if there exist two partitions

$$(\mathcal{R}_1^{\mathsf{V}}, \mathcal{R}_2^{\mathsf{V}}, \ldots, \mathcal{R}_p^{\mathsf{V}}) \quad \text{and} \quad (\mathcal{R}_1^{\mathsf{H}}, \mathcal{R}_2^{\mathsf{H}}, \ldots, \mathcal{R}_p^{\mathsf{H}})$$

of $\mathcal{R}'$, with $p \leq \ell$, as well as reals ν_t, μ_t for $t = 1, 2, \ldots, p$, with the following properties satisfied for each $t \in \{1, 2, \ldots, p\}$:

- $\nu_t \leq g_i < \mu_t$ for each $R_i \in \mathcal{R}_t^{\mathsf{V}}$;
- $\nu_t \leq h_i < \mu_t$ for each $R_i \in \mathcal{R}_t^{\mathsf{H}}$;
- $\nu_t/\mu_{t-1} = 1/\delta\epsilon$ (except for $t = 1$) and $\mu_t/\nu_t = (1/\delta\epsilon)^{(1/\epsilon)-1}$; and
- $\nu_1 \leq 1$, $\mu_p \geq L$, and all numbers μ_t and μ_t apart from ν_1 are integers.

The partitions $(\mathcal{R}_t^{\mathsf{V}})_{t=1,\ldots,p}$ and $(\mathcal{R}_t^{\mathsf{H}})_{t=1,\ldots,p}$ are called the *vertical* and *horizontal levels*, respectively, whereas the parameters $(\nu_t)_{t=1,\ldots,p}$ and $(\mu_t)_{t=1,\ldots,p}$ are the *lower* and *upper limits* of the corresponding levels. Note that, maybe somehow counter-intuitively, vertical levels partition $\mathcal{R}'$ by width, while the horizontal levels partition $\mathcal{R}'$ by heights; the reason for this will become clear later on. We now prove that we can find a well-separated subfamily that loses only an $\mathcal{O}(\epsilon)$-fraction of OPT using a standard shifting technique (see, e.g., Hochbaum and Maas [18]). Essentially the same step is used in the PTAS of Adamaszek et al. [1] (see Lemma 6 therein).

Lemma 1. *In polynomial time one can sample a subfamily $\mathcal{R}' \subseteq \mathcal{R}$ that is well-separated and satisfies $\mathsf{OPT}(\mathcal{R}') \geq (1 - 8\epsilon)\mathsf{OPT}$ with probability at least 3/4. Moreover, $\mathcal{R}'$ is constructed together with the corresponding partitions into vertical and horizontal levels, and with their lower and upper limits.*

Proof. Recall that the widths and heights of the rectangles from $\mathcal{R}$ are integers between 1 and $L-1$, where $L = (1/\delta\epsilon)^\ell$. Create first a partition of the rectangles into *vertical layers* $\mathcal{L}^{\mathsf{V}}_j$ for $j = 1, 2, \ldots, \ell$, where layer $\mathcal{L}^{\mathsf{V}}_j$ consists of rectangles R_i for which $(1/\delta\epsilon)^{j-1} \leq g_i < (1/\delta\epsilon)^j$. In a symmetric manner, partition $\mathcal{R}$ into *horizontal layers* $\mathcal{L}^{\mathsf{H}}_j$ for $j = 1, 2, \ldots, \ell$, where layer $\mathcal{L}^{\mathsf{H}}_j$ consists of rectangles R_i for which $(1/\delta\epsilon)^{j-1} \leq h_i < (1/\delta\epsilon)^j$.

Let us pick an offset b uniformly at random from the set $\{0, 1, \ldots, (1/\epsilon)-1\}$. Construct $\mathcal{R}'$ from $\mathcal{R}$ by removing all rectangles contained in those vertical layers $\mathcal{L}^{\mathsf{V}}_j$ and those horizontal layers $\mathcal{L}^{\mathsf{H}}_j$, for which $j \equiv b \mod (1/\epsilon)$. It is easy to see that $\mathcal{R}'$ constructed in this manner is well-separated: each vertical level $\mathcal{R}^{\mathsf{V}}_t$ consists of $(1/\epsilon)-1$ consecutive vertical layers between two removed ones, with the exception of the first and the last level, for which the start/end of the sequence of layers delimits the level. A symmetric analysis yields the partition into horizontal levels. Moreover, it is straightforward to compute in polynomial time the partition into horizontal/vertical levels, as well as to choose their lower and upper limits.

It remains to show that the probability that $\mathsf{OPT}(\mathcal{R}') \geq (1-8\epsilon)\mathcal{R}$ is at least 3/4. Fix any optimum solution $\mathcal{S}$ in $\mathcal{R}$, that is, an independent set of rectangles such that $w(\mathcal{S}) = \mathsf{OPT}$. Observe that for any rectangle $R_i \in \mathcal{S}$, the probability that the vertical layer it belongs to is removed during the construction of $\mathcal{R}'$, is equal to ϵ. Similarly, the probability that the horizontal layer to which R_i belongs is removed when constructing $\mathcal{R}'$, is also ϵ. Hence, R_i is not included in $\mathcal{R}'$ with probability at most 2ϵ. This means that the expected value of $w(\mathcal{S} \setminus \mathcal{R}')$, the total weight of rectangles from $\mathcal{S}$ that did not survive in $\mathcal{R}'$, is at most $2\epsilon \cdot \mathsf{OPT}$. By Markov's inequality, with probability at least 3/4 we have that $w(\mathcal{S} \setminus \mathcal{R}') \leq 8\epsilon \cdot \mathsf{OPT}$. This event, however, implies that $\mathsf{OPT}(\mathcal{R}') \geq w(\mathcal{S} \cap \mathcal{R}') \geq (1-8\epsilon)\mathsf{OPT}$, and hence we are done. □

From now on we work with the family $\mathcal{R}'$ obtained by Lemma 1, and we adopt the notation from Definition 1.

Hierarchical grid structure. Let $a \in \{1, \ldots, L-1\}$ be an integer shift parameter, to be determined later. Given a, we construct a hierarchy of grid lines in the plane. We start with horizontal lines, which will be divided into p levels corresponding to the horizontal levels $\mathcal{R}^{\mathsf{H}}_t$. For level t, define the *level-t unit* as $u_t = \delta\nu_t/2$. Note that for $t > 1$ we have $u_t = \mu_{t-1}/(2\epsilon)$, and hence u_t is an integer for $t > 1$, since $1/\epsilon$ is even. For each level $t \in \{1, 2, \ldots, p\}$ we define a set of horizontal grid lines G^{H}_t, consisting of the horizontal lines with y-coordinates from the set

$$\{a + b \cdot u_t \colon b \in \mathbb{Z}\}.$$

In other words, we take horizontal lines that are u_t apart from each other, and we shift them so that there is a line with y-coordinate a. We define vertical grid lines G^{V}_t of levels $t = 1, 2, \ldots, p$ in a symmetric manner, using the same shift parameter a and the same units for all levels. Define the *grid* of level t to be $G_t = G^{\mathsf{H}}_t \cup G^{\mathsf{V}}_t$. Note that any line of G_t is also a line of $G_{t'}$ for all $t' < t$. Thus, the grid of each level t' refines the grid of each larger level t.

Before we proceed, we briefly describe the intuition of the next step; this step is also present in the PTAS of Adamaszek et al. [1] (see Lemma 7 therein). Rectangles belonging to $\mathcal{R}^{\mathsf{V}}_{t'}$ for $t' \geq t$ have width not smaller than ν_t. On the other hand, the lines of G^{V}_t are spaced at distance $u_t = \delta\nu_t/2$ apart, which means that there are $\Omega(1/\delta)$ vertical grid lines of G^{V}_t crossing each rectangle of vertical level t or larger. Intuitively, G^{V}_t provides a fine grid for those vertical levels, so that their rectangles can be snapped to the lines of G^{V}_t via shrinking by a multiplicative factor of at most $1-\delta$. On the other hand, the rectangles of vertical levels $t-1$ or smaller have widths not larger than $\mu_{t+1} = \nu_t \cdot (\delta\epsilon)$. This means that the grid lines of G^{V}_t are actually at much larger distance from each other than the maximum possible

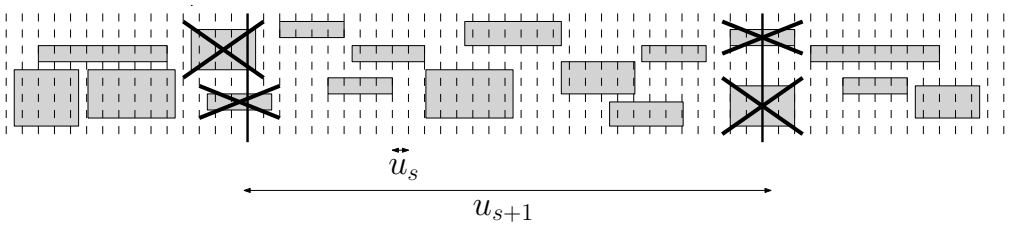

Figure 1: The vertical grid. The dashed vertical lines are the vertical grid lines of G_s^{V}, the bold vertical lines are the lines in the set G_{s+1}^{V}. All shown rectangles are from level $\mathcal{R}_s^{\mathsf{V}}$. The crossed out rectangles are abusive since they intersect lines from G_{s+1}^{V}.

width of such rectangles; more precisely, larger by a multiplicative factor at least $1/(2\epsilon)$. Consequently, if we choose the shift parameter $a \in \{1, \ldots, L-1\}$ uniformly at random, the probability that a rectangle R_i will be crossed by a vertical line of level larger than its vertical level, or a horizontal line of level larger than its horizontal level, will be $\mathcal{O}(\epsilon)$. If we exclude such rectangles, then we lose only an $\mathcal{O}(\epsilon)$-fraction of OPT in expectation, while achieving the property that vertical lines of each level t separate rectangles of lower vertical levels, and the symmetric claim holds for horizontal lines as well.

We now formalize the above intuition. Take a rectangle $R_i \in \mathcal{R}'$, and suppose that $R_i \in \mathcal{R}_s^{\mathsf{V}}$ and $R_i \in \mathcal{R}_t^{\mathsf{H}}$. We say that R_i is *abusive* if R_i is crossed by a vertical line of level larger than s, or R_i is crossed by a horizontal line of level larger than t; see Figure 1.

Lemma 2. *Suppose that a is sampled uniformly at random from $\{0, \ldots, L-1\}$. Then, for each $R_i \in \mathcal{R}'$, the probability that R_i is abusive is at most 8ϵ.*

Proof. Suppose $R_i \in \mathcal{R}_s^{\mathsf{V}} \cap \mathcal{R}_t^{\mathsf{H}}$. Rectangle R_i is crossed by a vertical line of level larger than s if and only if it is crossed by a vertical line of level $s+1$. Since $R_i \in \mathcal{R}_s^{\mathsf{V}}$, we have that

$$g_i \leq \mu_s \leq \nu_{s+1} \cdot (\delta\epsilon) = u_{s+1} \cdot 2\epsilon.$$

Lines of G_{s+1}^{V} are spaced at distance u_{s+1} from each other, which means that R_i is crossed by a line of G_{s+1}^{V} if and only if the remainder of a modulo u_{s+1} is among a set of $g_i - 1$ values, corresponding to the vertical lines with integral x-coordinates that cross R_i. As $u_{s+1} < L$ and a is drawn uniformly at random from $\{0, 1, \ldots, L-1\}$, this happens with probability at most $2 \cdot \frac{g_i - 1}{u_{s+1}} \leq 4\epsilon$. Symmetrically, R_i is crossed by a horizontal line of level larger than t with probability at most 4ϵ. Therefore, R_i is abusive with probability at most 8ϵ. □

We now construct a subfamily $\mathcal{R}'' \subseteq \mathcal{R}'$ by drawing a uniformly at random from the set $\{0, 1, \ldots, L-1\}$, defining the grid lines $(G_t^{\mathsf{V}})_{t=1,\ldots,p}$ and $(G_t^{\mathsf{H}})_{t=1,\ldots,p}$, and removing all abusive rectangles from $\mathcal{R}'$. Lemma 2 implies the following.

Lemma 3. $\mathsf{OPT}(\mathcal{R}'') \geq (1 - 32\epsilon)\mathsf{OPT}(\mathcal{R}')$ *with probability at least* $3/4$.

Proof. Let $\mathcal{S} \subseteq \mathcal{R}'$ be such that $w(\mathcal{S}) = \mathsf{OPT}(\mathcal{R}')$. Since each rectangle of $\mathcal{R}'$ is abusive with probability at most 8ϵ, we have that the expected total weight of rectangles of $\mathcal{S}$ that turns out to be abusive is at most $8\epsilon \cdot \mathsf{OPT}(\mathcal{R}')$. By Markov's inequality, the probability that this total weight exceeds $32\epsilon \cdot \mathsf{OPT}(\mathcal{R}')$ is at most $1/4$. Consequently, with probability at least $3/4$ we have that $\mathsf{OPT}(\mathcal{R}'') \geq w(\mathcal{S} \cap \mathcal{R}'') \geq (1 - 32\epsilon)\mathsf{OPT}(\mathcal{R}')$. □

Since the sampling of a is done independently of the sampling used for the construction of $\mathcal{R}'$, Lemmas 1 and 3 combined imply that

$$\mathsf{OPT}(\mathcal{R}'') \geq (1 - 32\epsilon)(1 - 8\epsilon)\mathsf{OPT}(\mathcal{R}) \geq (1 - 40\epsilon)\mathsf{OPT}(\mathcal{R}), \tag{1}$$

with probability at least $3/4 \cdot 3/4 > 1/2$. Moreover, $\mathcal{R}''$ is still well-separated, where the horizontal and vertical levels are $\mathcal{R}^{\mathsf{H}}_t$ and $\mathcal{R}^{\mathsf{V}}_t$ trimmed to $\mathcal{R}''$, and their limits are unchanged. Finally, $\mathcal{R}''$ does not contain any abusive rectangle.

Snapping by shrinking. When considering $\mathcal{R}''$, the vertical grid lines of G^{V}_t provide a fine division of every rectangle from vertical level t or larger, while no rectangle of smaller vertical level is crossed by them; a symmetric claim holds also for horizontal grid lines. The idea now is to shrink each rectangle $R_i \in \mathcal{R}''$ so that its vertical sides are aligned with some vertical grid lines of the vertical level of R_i, while the horizontal sides are aligned with some horizontal grid lines of the horizontal level of R_i. This is formalized in the next lemma, which is also similar to Adamaszek et al. [1]. However, in this step there is a subtle but crucial difference. Consider a rectangle $R_i \in \mathcal{R}''$, and suppose that $R_i \in \mathcal{R}^{\mathsf{V}}_s$ and $R_i \in \mathcal{R}^{\mathsf{H}}_t$. In [1] the rectangle R_i is shrunk in the vertical dimension only if $s \geq t$ and in the horizontal dimension only if $t \geq s$. In this paper we always do both which will be important for our dynamic programming algorithm later.

Lemma 4. *In polynomial time we can compute a family of axis-parallel rectangles $\mathcal{Q}$ that contains one rectangle Q_i for each $R_i \in \mathcal{R}''$, of the same weight w_i as R_i and satisfying the following conditions:*

- *$R_i^{-\delta} \subseteq Q_i \subseteq R_i$ for each $R_i \in \mathcal{R}''$; and*
- *if $R_i \in \mathcal{R}^{\mathsf{V}}_s \cap \mathcal{R}^{\mathsf{H}}_t$, then both vertical sides of Q_i are contained in some vertical grid lines of G^{V}_s, and both horizontal sides of Q_i are contained in some horizontal grid lines of G^{H}_t.*

Proof. Take any $R_i \in \mathcal{R}''$, and suppose $R_i \in \mathcal{R}^{\mathsf{V}}_s \cap \mathcal{R}^{\mathsf{H}}_t$. We define Q_i as the rectangle cut from the plane by the following four lines:

- the left-most and the right-most vertical grid lines of G^{V}_s that cross R_i;
- the bottom-most and the top-most horizontal grid lines of G^{H}_t that cross R_i.

Clearly, we have that $Q_i \subseteq R_i$ and the second condition of the statement is satisfied. We are left with proving that $R_i^{-\delta} \subseteq Q_i$.

Consider first the left side of Q_i, which is contained in the left-most vertical grid line of G^{V}_s that crosses R_i. Since $R_i \in \mathcal{R}^{\mathsf{V}}_s$, we have that $g_i \geq \nu_s$, while the grid lines of G^{V}_s are spaced at distance $u_t = \delta\nu_s/2$ apart. This means that the left-most vertical grid line crossing R_i has the x-coordinate not larger than $x_i^{(1)} + \delta\nu_s/2$, which in turn is not larger than $x_i^{(1)} + \delta g_i/2$. This means that the left side of Q_i is either to the left or at the same x-coordinate as the left side of $R_i^{-\delta}$. An analogous reasoning can be applied to the other three sides of Q_i, thereby proving that $R_i^{-\delta} \subseteq Q_i$. □

In the next subsection we will actually show that we can find an optimum solution for the family $\mathcal{Q}$ using dynamic programming within the claimed running time. By (1) and the first condition of Lemma 4, we have

$$\mathsf{OPT}(\mathcal{Q}) \geq \mathsf{OPT}(\mathcal{R}'') \geq (1 - 40\epsilon)\mathsf{OPT}$$

with probability at least $1/2$. Hence, the optimum solution for $\mathcal{Q}$ indeed has large enough weight with good probability. Moreover, by the first condition of Lemma 4, for any independent set of rectangles in $\mathcal{Q}$, the corresponding rectangles in $\mathcal{R}^{-\delta}$ are also independent. Hence, any solution for MWISR on $\mathcal{Q}$ projects to a solution of the same weight for MWISR with δ-shrinking on $\mathcal{R}$.

From now on we focus on the family $\mathcal{Q}$. For each $t \in \{1, 2, \ldots, p\}$, let

$$\mathcal{Q}^{\mathsf{V}}_t = \{Q_i \colon R_i \in \mathcal{R}^{\mathsf{V}}_t\} \quad \text{and} \quad \mathcal{Q}^{\mathsf{H}}_t = \{Q_i \colon R_i \in \mathcal{R}^{\mathsf{H}}_t\}.$$

Note that since $\mathcal{Q}$ is obtained only by shrinking rectangles from $\mathcal{R}''$, it is still the case that no rectangle of $\mathcal{Q}$ is abusive.

3.2 Dynamic programming

We now present a dynamic programming algorithm that, given the family $\mathcal{Q}$ constructed in the previous section, computes the value $\mathsf{OPT}(\mathcal{Q})$. An optimum solution, that is, an independent set of rectangles of total weight $\mathsf{OPT}(\mathcal{Q})$, can be recovered from the run of the dynamic program using standard methods, and hence for simplicity we omit this aspect in the description.

Actually, it will be convenient to describe the algorithm as *backtracking with memoization*. That is, the subproblems are solved by recursion, but once a subproblem has been solved once, the optimum value for it is stored in a map (is *memoized*), and further calls to solving this subproblem will only retrieve the memoized optimum value, rather than solve the subproblem again. Solving each subproblem (excluding the recursive subcalls) will take time $f(\delta,\epsilon)\cdot n^{\mathcal{O}(1)}$ for some computable function f, and we will argue that at most $g(\delta,\epsilon)\cdot(nN)^{\mathcal{O}(1)}$ subproblems are solved in total, for some other computable function g. This will ensure the promised running time of the algorithm.

We first define subproblems. A *subproblem* is a tuple $I=(s,t,x_1,x_2,y_1,y_2)$:

- the pair $(s,t)\in\{1,\ldots,p\}\times\{1,\ldots,p\}$ is the *level* of the subproblem, which consists of the vertical level s and the horizontal level t; and
- x_1,x_2,y_1,y_2 are integers satisfying

$$x_1 < x_2 \leq x_1 + (1/\delta\epsilon)^{1/\epsilon} \quad \text{and} \quad y_1 < y_2 \leq y_1 + (1/\delta\epsilon)^{1/\epsilon}.$$

  Integers x_1,x_2 are the lower and upper *vertical offsets*, respectively, while y_1,y_2 are the lower and upper *horizontal offsets*, respectively.

The *area covered* by subproblem $I=(s,t,x_1,x_2,y_1,y_2)$ is the rectangle

$$A_I = (a+x_1\cdot u_s, a+x_2\cdot u_s)\times(a+y_1\cdot u_t, a+y_2\cdot u_t).$$

In other words, (x_1,x_2,y_1,y_2) define the offsets of the four grid lines—two from G^{V}_s and two from G^{H}_t—that cut out A_I from the plane.

For a subproblem I, let $\mathcal{Q}_I$ be the set of all rectangles from $\mathcal{Q}$ that are contained in A_I. The next check follows from a simple calculation of parameters.

Lemma 5. *If I is a subproblem of level (s,t), then $\mathcal{Q}_I \subseteq \bigcup_{s'\leq s,\, t'\leq t} \mathcal{Q}^{\mathsf{V}}_{s'} \cap \mathcal{Q}^{\mathsf{H}}_{t'}$.*

Proof. Consider a rectangle Q_i that belongs to a vertical level $\mathcal{Q}^{\mathsf{V}}_{s'}$ for some $s'>s$. By the first condition of Lemma 4, we have that the width of Q_i is at least $(1-\delta)g_i$. However, since $s'>s$, we have $g_i \geq \nu_{s'} \geq \nu_s\cdot(1/\delta\epsilon)^{1/\epsilon}$. On the other hand, the width of A_I is at most $(1/\delta\epsilon)^{1/\epsilon}\cdot u_s = \nu_s\cdot(1/\delta\epsilon)^{1/\epsilon}\cdot(\delta/2)$. Since we assumed that $1/\delta>4$, we infer that the width of Q_i is larger than the width of A_I, and hence Q_i cannot be contained in A_I. A symmetric reasoning shows also that any rectangle that belongs to a horizontal level $\mathcal{Q}^{\mathsf{H}}_{t'}$ for some $t'>t$ cannot belong to $\mathcal{Q}_I$ due to having larger height than A_I. The claim follows. □

For a subproblem I, we define the *value* of I, denoted $\mathsf{Value}(I)$, as follows:

$$\mathsf{Value}(I) = \max\{\, w(\mathcal{S}) \,:\, \mathcal{S}\subseteq\mathcal{Q}_I \text{ and } \mathcal{S} \text{ is independent}\,\}.$$

We first note that there is a subproblem that encompasses the whole instance.

Lemma 6. *There is a subproblem I_{all} of level (p,p), computable in constant time, such that $A_{I_{\mathsf{all}}} \supseteq (1,L)\times(1,L)$. Consequently, $\mathsf{OPT}(\mathcal{Q}) = \mathsf{Value}(I_{\mathsf{all}})$.*

Proof. Choose z_1 to be the largest integer such that $a + z_1 \cdot u_p \leq 0$. Let $z_2 = z_1 + (1/\delta\epsilon)^{1/\epsilon}$, and define $I_{\mathsf{all}} = (p, p, z_1, z_1, z_2, z_2)$. Clearly I_{all} is a subproblem, so we are left with proving that $A_{I_{\mathsf{all}}} \supseteq (1, L) \times (1, L)$. By the maximality of z_1,

$$a + z_1 \cdot u_p > -u_p. \tag{2}$$

On the other hand, we have

$$u_p = \delta\nu_p/2 = \mu_p \cdot (\delta/2) \cdot (\delta\epsilon)^{1/\epsilon - 1} \geq L \cdot \delta/2 \cdot (\delta\epsilon)^{1/\epsilon - 1}. \tag{3}$$

By combining (2) and (3), we obtain

$$\begin{aligned} a + z_2 \cdot u_p &= a + z_1 \cdot u_p + (1/\delta\epsilon)^{1/\epsilon} \cdot u_p \geq -u_p + (1/\delta\epsilon)^{1/\epsilon} \cdot u_p \\ &\geq L \cdot \delta/2 \cdot (\delta\epsilon)^{1/\epsilon - 1} \cdot ((1/\delta\epsilon)^{1/\epsilon} - 1) > L; \end{aligned}$$

the last inequality follows from the assumption that $1/\epsilon$ is an integer larger than 4. It follows that $A_{I_{\mathsf{all}}}$ indeed contains the whole square $(1, L) \times (1, L)$. □

Next, we show how to *solve* each subproblem I, that is, to compute $\mathsf{Value}(I)$, using recursion.

Lemma 7. *A subproblem I of level (s, t) can be solved using $f(\delta, \epsilon)$ calls to solving subproblems of levels $(s-1, t)$, $(s, t-1)$, and $(s-1, t-1)$, for some computable function f. Moreover, the time needed for this computation, excluding the time spent in the recursive calls, is at most $f(\delta, \epsilon) \cdot n$.*

Proof. Let $I = (s, t, x_1, x_2, y_1, y_2)$. We denote $x = x_2 - x_1$ and $y = y_2 - y_1$; recall that $0 < x, y \leq (1/\delta\epsilon)^{1/\epsilon}$. Consider all vertical grid lines of G_s^{V} with x-coordinates $a + \alpha \cdot u_s$ for $x_1 \leq \alpha \leq x_2$ and all horizontal grid lines of G_t^{H} with y-coordinates $a + \beta \cdot u_t$ for $y_1 \leq \beta \leq y_2$. These lines partition the rectangle A_I into $x \times y$ smaller rectangles with side lengths u_s and u_t, which we shall call *cells*. Cells can be naturally indexed by pairs of integers from $\{1, \ldots, x\} \times \{1, \ldots, y\}$: cell (α, β) has sides contained in vertical lines with x-coordinates $a + (\alpha - 1) \cdot u_s$ and $a + \alpha \cdot u_s$, and horizontal lines with y-coordinates $a + (\beta - 1) \cdot u_t$ and $a + \beta \cdot u_t$. We use the terminology of *rows* and *columns* in the grid of cells, defined naturally.

Consider any independent set of rectangles $\mathcal{S} \subseteq \mathcal{Q}_I$. The reader may think of $\mathcal{S}$ as an optimum solution, that is, a set with $w(\mathcal{S}) = \mathsf{Value}(I)$. By Lemma 6, each rectangle of $\mathcal{S}$ belongs to $\mathcal{Q}_{s'}^{\mathsf{V}} \cap \mathcal{Q}_{t'}^{\mathsf{H}}$ for some $s' \leq s$ and $t' \leq t$; such rectangle will be henceforth called a rectangle of *level* (s', t'). Let us classify these rectangles according to their level as follows. A rectangle $Q_i \in \mathcal{S}$ of level (s', t') is called:

- *large* if $s' = s$ and $t' = t$;
- *horizontal* if $s' = s$ and $t' < t$;
- *vertical* if $s' < s$ and $t' = t$;
- *small* if $s' < s$ and $t' < t$.

Observe that if $s = 1$ then there are no vertical or small rectangles, whereas if $t = 1$ then there are no horizontal or small rectangles. The following assertions follow from the second condition of Lemma 4, and the fact that no rectangle of $\mathcal{Q}$ is abusive; see Figure 2 for an illustration.

- Each large rectangle is equal to the union of a subset of cells.
- Each horizontal rectangle is contained in the union of a sequence of consecutive cells in a single row of the grid. The vertical sides of this rectangle are aligned with the left side of the first cell of the sequence, and the right side of the last cell.

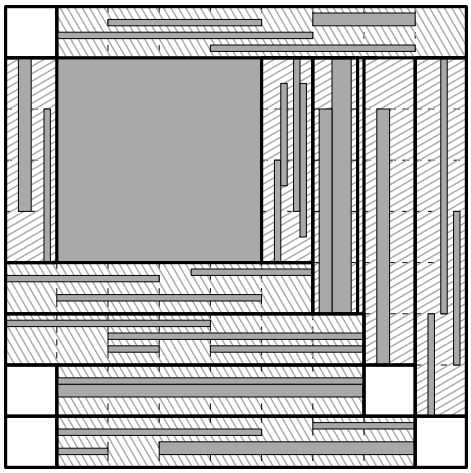

Figure 2: The partition of A_I into large, horizontal, vertical, and small cells. The figure is a slightly modified figure from [1].

- Each vertical rectangle is contained in the union of a sequence of consecutive cells in a single column of the grid. The horizontal sides of this rectangle are aligned with the bottom side of the first cell of the sequence, and the top side of the last cell.
- Each small rectangle is contained in a single cell.

We shall call a cell:

- *large* if it is contained in a large rectangle;
- *horizontal* if it overlaps with some horizontal rectangle;
- *vertical* if it overlaps with some vertical rectangle; and
- *small* if all rectangles from $\mathcal{S}$ with which it overlaps are small.

Observe that as $\mathcal{S}$ is independent, each cell is of exactly one of the above types.

We now partition the cells into *boxes*, where each box is formed by a rectangle of cells; again, see Figure 2 for an illustration. First, for each large rectangle, create a *large box* consisting of the cells contained in this rectangle. Second, for every inclusion-wise maximal sequence of consecutive horizontal cells contained in the same row, create a *horizontal box* consisting of all these cells. Third, for every inclusion-wise maximal sequence of consecutive vertical cells contained in the same column, create a *vertical box* consisting of all these cells. Finally, put every small cell into a *small box* consisting only of this cell. Observe that thus, every rectangle of $\mathcal{S}$ is contained in a single box.

Suppose for a moment that $s > 1$ and consider any potential vertical box B, that is, a sequence of consecutive cells contained in one column of the grid of cells. Observe that for this box we can define a subproblem I_B of level $(s-1, t)$ so that $B = A_{I_B}$. This is because the grid lines delimiting B are also grid lines from G^{V}_{s-1} and G^{H}_t, while the width of B is equal to $u_s = (1/\delta\epsilon)^{1/\epsilon} u_{s-1}$. Hence, for the subproblem I_B the difference between its vertical offsets will be bounded by $(1/\delta\epsilon)^{1/\epsilon}$, as requested in the definition of a subproblem. Similarly we can define a subproblem I_B of level $(s, t-1)$ whenever B is a horizontal box, and a subproblem I_B of level $(s-1, t-1)$ whenever B is a small box.

With these observations in mind, we can present the procedure for computing $\mathsf{Value}(I)$; suppose for now that $s, t > 1$. The procedure enumerates all partitions of the grid of cells into boxes, where

each box has a prescribed type: it is either large, horizontal, vertical, or small. Obviously, we consider only partitions where horizontal boxes are contained in one row, vertical boxes are contained in one column, and small boxes consist of one cell. Observe that since $x, y \leq (1/\delta\epsilon)^{1/\epsilon}$, the number of such partitions is bounded by $f(\delta, \epsilon)$ for some computable function f. For each partition $\mathcal{P}$, we do the following:

- For each large box B in $\mathcal{P}$, choose the heaviest rectangle $Q_i \in \mathcal{Q}_I$ such that $B = Q_i$, and record its weight. If for some large box there is no such rectangle, ignore this partition $\mathcal{P}$.
- For each horizontal, vertical, or small box B in $\mathcal{P}$, compute $\mathsf{Value}(I_B)$ by a recursive call to solving the subproblem I_B.
- Sum up the weights recorded for large boxes and values computed for other boxes, and record the obtained sum as a candidate value for $\mathsf{Value}(I)$.

Finally, return as $\mathsf{Value}(I)$ the maximum among all the candidate values computed for all partitions into boxes. In case $s = 1$ we assume that all vertical and small boxes yield value 0, and when $t = 1$ we do the same for all horizontal and small boxes. The claimed running time of the procedure, as well as the upper bound on the number of recursive calls, follows directly from the description.

To see the correctness of the procedure, observe that each candidate value represents the total weight of some independent set of rectangles contained in $\mathcal{Q}_I$: this independent set consists of all the rectangles chosen to fill the large boxes, and the union of the independent sets yielding values $\mathsf{Value}(I_B)$ among all the other boxes B. On the other hand, from the discussion before the description of the procedure it follows that if $\mathcal{S}_I \subseteq \mathcal{Q}_I$ is an independent set of rectangles such that $w(\mathcal{S}_I) = \mathsf{Value}(I)$, then there will be a partition $\mathcal{P}$ into boxes which will be considered by the procedure, and for which every rectangle of $\mathcal{S}_I$ fits into a single box. It is then easy to see that for this partition $\mathcal{P}$, the optimum value $\mathsf{Value}(I)$ will be computed as a candidate. □

Finally, to bound the running time of the algorithm, we prove that there is only a small number of subproblems I for which $\mathcal{Q}_I$ is non-empty. Obviously, only such subproblems are necessary to solve, as the others have value 0.

Lemma 8. *The number of subproblems I for which $\mathcal{Q}_I$ is non-empty is at most $81 \cdot (1/\delta\epsilon)^{4/\epsilon} \cdot |\mathcal{Q}|p^2$.*

Proof. It suffices to prove that for every fixed rectangle $Q_i \in \mathcal{Q}$ and for every level $(s, t) \in \{1, \ldots, p\} \times \{1, \ldots, p\}$, there are at most $81 \cdot (1/\delta\epsilon)^{4/\epsilon}$ subproblems I of level (s, t) for which $Q_i \in \mathcal{Q}_I$. Observe that if $I = (s, t, x_1, x_2, y_1, y_2)$ and $I' = (s, t, x'_1, x'_2, y'_1, y'_2)$ are two subproblems of level (s, t) such that $A_I \cap A_{I'} \neq \emptyset$, then we necessarily have

$$x_1 - (1/\delta\epsilon)^{1/\epsilon} < x'_1 < x'_2 < x_2 + (1/\delta\epsilon)^{1/\epsilon} \quad \text{and}$$
$$y_1 - (1/\delta\epsilon)^{1/\epsilon} < y'_1 < y'_2 < y_2 + (1/\delta\epsilon)^{1/\epsilon}.$$

Since $x_2 \leq x_1 + (1/\delta\epsilon)^{1/\epsilon}$ and $y_2 \leq y_1 + (1/\delta\epsilon)^{1/\epsilon}$, for fixed I there are at most $81 \cdot (1/\delta\epsilon)^{4/\epsilon}$ choices of x'_1, x'_2, y'_1, y'_2 that satisfy the above. Therefore, for each subproblem I of level (s, t), there are at most $81 \cdot (1/\delta\epsilon)^{4/\epsilon}$ other subproblems I' of the same level for which $A_I \cap A_{I'} \neq \emptyset$. If $Q_i \in \mathcal{Q}_I$ for some I of level (s, t), then all subproblems I' of the same level for which $Q_i \in \mathcal{Q}_{I'}$ must satisfy $A_I \cap A_{I'} \neq \emptyset$, and hence there can be at most $81 \cdot (1/\delta\epsilon)^{4/\epsilon}$ of them. □

Having gathered all the tools, we conclude by describing the algorithm. To compute $\mathsf{OPT}(\mathcal{Q})$, it is sufficient to compute $\mathsf{Value}(I_{\mathsf{all}})$ for the subproblem I_{all} given by Lemma 6. As mentioned before, for this we use backtracking with memoization. Namely, starting from I_{all}, we recursively solve subproblems as explained in Lemma 7. Whenever $\mathsf{Value}(I)$ has been computed for some

subproblem I, then this value is memoized in a map, and further calls to solving I will only return the value retrieved from the map, instead of recomputing the value again. Furthermore, whenever we attempt to compute $\mathsf{Value}(I)$ for a subproblem I for which $\mathcal{Q}_I$ is empty, we immediately return 0 instead of applying the procedure of Lemma 7. Therefore, the total running time of the algorithm is upper bounded by the number of subproblems I for which $\mathcal{Q}_I$ is non-empty, times the time spent on internal computations for each of them, including checking whether the respective set $\mathcal{Q}_I$ is empty and whether $\mathsf{Value}(I)$ has already been memoized. The first factor is bounded by $81 \cdot (1/\delta\epsilon)^{4/\epsilon} \cdot |\mathcal{Q}|p^2 \leq 81 \cdot (1/\delta\epsilon)^{4/\epsilon} \cdot nN^2$ due to Lemma 8. The second factor is bounded by $f(\delta,\epsilon) \cdot n^d$ for some constant d, due to Lemma 7. Hence, the running time of the whole algorithm is $f(\delta,\epsilon) \cdot (nN)^c$ for some computable function f and constant c. As mentioned before, the algorithm can be trivially adjusted to also compute an independent set of weight $\mathsf{Value}(I_{\mathsf{all}})$ by storing the value of each subproblem together with some independent set that certifies this value.

Summarizing, the dynamic programming described above computes an independent set in $\mathcal{Q}$ of weight $\mathsf{OPT}(\mathcal{Q})$ in time $f(\delta,\epsilon) \cdot (nN)^c$. As argued in the previous section, such an independent set projects to an independent set of the same weight in $\mathcal{R}^{-\delta}$, and $\mathsf{OPT}(\mathcal{Q}) \geq (1-40\epsilon)\mathsf{OPT}$ holds with probability at least $1/2$. This concludes the proof of Theorem 2.

3.3 Parameterized setting: Theorem 1

In this section we present the proof of Theorem 1. Note that if all rectangles had equal weights, then we could derive Theorem 1 from Theorem 2 by simply running the algorithm for $\epsilon = 1/(k+1)$. This is because in the uniform-weight setting, a solution of size at most k that loses at most a multiplicative factor of $1/(k+1)$ to the optimum, actually needs to be optimal. However, this argument breaks when the weights are non-uniform. Instead, we argue how the reasoning from the proof of Theorem 2 should be adjusted to also justify Theorem 1.

First, we consider the sparsification arguments presented in Section 3.1, which yield subfamilies $\mathcal{R}'$ and $\mathcal{R}''$. Let us apply them for $\epsilon = 1/(20k)$. Suppose $\mathcal{S} \subseteq \mathcal{R}$ is an optimum solution for the considered variant of the problem: it is an independent set of size at most k and of total weight $\mathsf{OPT}_k(\mathcal{R})$. When applying the sparsification yielding family $\mathcal{R}'$ (Lemma 1), every rectangle of $\mathcal{R}$ is removed with probability at most $2\epsilon = 1/(10k)$ (cf. the proof of Lemma 1). Therefore, the expected number of rectangles from $\mathcal{S}$ that are removed when constructing $\mathcal{R}'$ is at most $1/10$. Hence, by Markov's inequality, with probability at least $9/10$ none of them is removed, and $\mathcal{S}$ entirely survives in $\mathcal{R}'$. Similarly, when $\mathcal{R}''$ is constructed from $\mathcal{R}'$, every rectangles is again removed with probability at most $8\epsilon = 2/(5k)$ (Lemma 2). Again, the same calculation using Markov's inequality yields that, conditioned on $\mathcal{S}$ surviving in $\mathcal{R}'$, $\mathcal{S}$ survives also in $\mathcal{R}''$ with probability at least $3/5$. In total, we conclude that $\mathcal{S}$ survives in $\mathcal{R}''$ with probability at least $9/10 \cdot 3/5 > 1/2$. Note that if $\mathcal{S} \subseteq \mathcal{R}''$, then $\mathsf{OPT}_k(\mathcal{R}'') \geq w(\mathcal{S}) = \mathsf{OPT}_k(\mathcal{R})$, but also $\mathsf{OPT}_k(\mathcal{R}'') \leq \mathsf{OPT}_k(\mathcal{R})$ due to $\mathcal{R}'' \subseteq \mathcal{R}$. Consequently, in this case we have $\mathsf{OPT}_k(\mathcal{R}'') = \mathsf{OPT}_k(\mathcal{R})$.

The construction of $\mathcal{Q}$ from $\mathcal{R}''$ is kept unchanged. Finally, in the dynamic programming we would like to compute $\mathsf{OPT}_k(\mathcal{Q})$ instead of $\mathsf{OPT}(\mathcal{Q})$. The adjustment is standard: we add one additional coordinate to the definition of a subproblem, which keeps track of the number of rectangles to be picked in the solution. More precisely, a subproblems is now a tuple $I = (s,t,x_1,x_2,y_1,y_2,\lambda)$, where the first six coordinates have the original meaning, while λ is an integer with $0 \leq \lambda \leq k$ that denotes the cardinality of the sought solution; λ will be called the *budget*. In subproblem I we are asked to compute $\mathsf{Value}(I)$, defined to be the largest weight of an independent set of size at most λ contained in $\mathcal{Q}_I$. The recursive formula for computing $\mathsf{Value}(I)$, presented in Lemma 7, needs to be adjusted as follows: for each of the considered partitions into boxes, we also consider all possible partitions of the budget among the boxes. Each large box has to receive budget 1, since

there will be exactly one rectangle fit into this box, and the budget λ_B allocated to any other box B becomes the budget in the corresponding subproblem I_B that is solved recursively. Note that since the number of boxes is bounded by $(1/\delta\epsilon)^{4/\epsilon}$ where $\epsilon = 1/(20k)$, we have that the number of partitions of the budget λ among the boxes is bounded by a computable function of δ and k.

Wrapping up, with probability at least $1/2$ we have that $\mathsf{OPT}_k(\mathcal{R}'') = \mathsf{OPT}_k(\mathcal{R})$. Since $\mathcal{Q}$ is obtained from $\mathcal{R}''$ only by shrinking the rectangles, we have $\mathsf{OPT}_k(\mathcal{Q}) \geq \mathsf{OPT}_k(\mathcal{R}'')$. As before, any independent set of rectangles in $\mathcal{Q}$ projects to an independent set of the same weight in $\mathcal{R}^{-\delta}$. Hence, if indeed $\mathsf{OPT}_k(\mathcal{R}'') = \mathsf{OPT}_k(\mathcal{R})$, then the dynamic programming will output a set of rectangles of size at most k that (after projection) is independent in $\mathcal{R}^{-\delta}$ and has weight at least $\mathsf{OPT}_k(\mathcal{R})$.

4 Kernelization results

In this section we present our kernelization results, which apply to the case when $\mathcal{R}$ consists of squares. As mentioned in the introduction, we can obtain a polynomial kernel for three settings: unit squares of uniform weight, unit squares of non-uniform weight, and squares of non-uniform size and uniform weight.

We start our description by clarifying the definition of kernelization. Then, we present two auxiliary lemmas on (weighted) graphs. Later, we apply these lemmas to reduced instances of MWISR on (unit) squares to obtain the different kernels. Finally, we explain how to reduce the bit encodings of the kernels, and we present applications for the design of EPTASes and fast FPT algorithms.

4.1 Definition of kernel

The classic definition of kernelization (cf. [11]) is tailored to decision problems; extending it to optimization problems in a weighted setting often turns out to be problematic. Making the definition compatible with the δ-shrinking model would make it even more complicated. For this reason, we define explicitly what we mean by kernelization for MWISR in the shrinking model, bearing in mind the main principle of kernelization: solving the kernel should project to a solution for the original instance.

Definition 2. Suppose $\mathcal{R}$ is a family of axis-parallel rectangles, k is an integer parameter, and $\delta \in (0,1)$ is a constant. Then, a *kernel for $\mathcal{R}$ and δ* is a polynomial-time computable subfamily $\mathcal{Q} \subseteq \mathcal{R}$ such that $|\mathcal{Q}| \leq f(k,\delta)$ for a computable function f, called the *size* of the kernel, and the following holds:

$$\mathsf{OPT}_k(\mathcal{Q}^{-\delta}) \geq \mathsf{OPT}_k(\mathcal{R}).$$

Thus, solving MWISR on $\mathcal{R}$ with δ-shrinking relaxation can be done by solving MWISR on $\mathcal{Q}^{-\delta}$ (without δ-shrinking relaxation). If one wishes to use an algorithm for the shrinking relaxation on the kernel, then applying it to $\mathcal{Q}^{-\delta}$ with parameter δ will yield a subfamily $\mathcal{S}$ of size k such that $\mathcal{S}^{-2\delta}$ is independent, and $w(\mathcal{S}) \geq \mathsf{OPT}_k(\mathcal{Q}^{-\delta}) \geq \mathsf{OPT}_k(\mathcal{R})$. Hence, this solves the original problem for 2δ-shrinking, and we can rescale the parameter accordingly.

A reader well-versed in the foundations of parameterized complexity will notice that Definition 2 is at least as strong as any reasonable complexity-theoretical definition of kernelization, apart from one aspect. Namely, the weights and the coordinates of the rectangles in the kernel are inherited from the original instance, and hence, their bit encoding may not be bounded in terms of k and δ. We prefer to work with Definition 2, because it focusses our efforts on the core combinatorial aspects of our kernelization procedures, rather than on the less relevant complexity-theoretical aspects. However, at the end of this section we argue that, essentially, the bit encodings of both the weights and the coordinates can be reduced to a polynomial in k and $1/\delta$.

4.2 Auxiliary results on graphs

We consider the parameterized Maximum Weight Independent Set (MWIS) problem on graphs. In this problem, we are given a graph G, a non-negative weight function $w\colon V(G) \to \mathbb{R}_{\geq 0}$, and an integer parameter k; we look for a subset of $V(G)$ that has size at most k, contains only pairwise non-adjacent vertices (i.e., it is an *independent set*), and has the maximum possible weight subject to these conditions. This maximum weight will be denoted by $\mathsf{OPT}_k(G)$. We present two simple reduction techniques that apply to parameterized MWIS on general graphs.

We need the following notation. Given a graph G, let $N_G^{\leq i}(v)$ denote the set of vertices at distance at most i from $v \in V(G)$, and let $\Delta^{\leq i}(G)$ denote the maximum number of vertices at distance at most i from any vertex of G (i.e., $\Delta^{\leq i}(G) = \max_{v \in V(G)} |N_G^{\leq i}(v)|$). Note that $\Delta^{\leq 1}(G)$ is simply the maximum degree of G plus 1.

In the next lemma, we will identify a subgraph G' of G such that it suffices to search in G' for the optimal solution to parameterized MWIS. Intuitively, later we will show that in the shrinking model for uniform squares the graph G' can be assumed to have size $64k/\delta^2$ (weighted case), since the shrinking will allow us to ensure that $\Delta^{\leq 2}(G) \leq 64/\delta^2$, for some appropriately chosen graph G.

Lemma 9. *Given an instance (G, w, k) of parameterized* MWIS*, one can in polynomial time find an induced subgraph G' of G such that $\mathsf{OPT}_k(G) = \mathsf{OPT}_k(G')$ and $|V(G')| \leq k \cdot \Delta^{\leq 2}(G)$.*

Proof. If G is empty, then the lemma holds with $G' = G$, and thus we may assume that G is not empty. We iteratively define a sequence of graphs $G_0, G_1, \ldots, G_k$ and a sequence of vertices $v_1, \ldots, v_k \in V(G)$ as follows. First, we put $G_0 = G$. Consider $i = 1, \ldots, k$ in order. If G_{i-1} is not empty, then let v_i be a vertex of maximum weight in G_{i-1}, and let $G_i = G_{i-1} - N_{G_{i-1}}^{\leq 2}[v_i]$. If G_{i-1} is empty, then let $v_i = v_{i-1}$ and $G_i = G_{i-1}$.

We claim there is a set $I \subseteq V(G)$, such that I is an independent set, $|I| \leq k$, I has maximum weight subject to these conditions, and $I \cap V(G_k) = \emptyset$. If G_k is an empty graph, then this is trivial. So assume that G_k is not empty. Then, $v_1, \ldots, v_k$ are all distinct by construction. Let $I' \subseteq V(G)$ be such that I' is an independent set, $|I'| \leq k$, I' has maximum weight under these conditions, and among all such sets $|I' \cap V(G_k)|$ is minimum. Suppose, for sake of contradiction, that $|I' \cap V(G_k)| > 0$. Let x be an arbitrary vertex in $I' \cap V(G_k)$. By construction, the sets $N_G^{\leq 1}[v_1], \ldots, N_G^{\leq 1}[v_k]$ are pairwise disjoint, because if $N_G^{\leq 1}[v_j] \cap N_G^{\leq 1}[v_i] \neq \emptyset$ for some $j > i$, then $v_j \in N_{G_{i-1}}^{\leq 2}[v_i]$, which combined with the fact that $v_j \neq v_i$ contradicts the choice of v_j. Moreover, $V(G_k)$ is disjoint from the sets $N_G^{\leq 1}[v_1], \ldots, N_G^{\leq 1}[v_k]$ as well. As $|I'| \leq k$ and $x \in I' \cap V(G_k)$, it follows that I' intersects at most $k-1$ sets among $N_G^{\leq 1}[v_1], \ldots, N_G^{\leq 1}[v_k]$. Consider any $i \in \{1, \ldots, k\}$ such that $I' \cap N_G^{\leq 1}[v_i] = \emptyset$. By the choice of $v_1, \ldots, v_k$, we have $w(v_i) \geq w(x)$. Therefore, the set $I'' = (I \setminus \{x\}) \cup \{v_i\}$ is an independent set, and it satisfies $|I''| = |I'|$, $w(I'') \geq w(I')$, and $|I'' \cap V(G_k)| < |I' \cap V(G_k)|$; this contradicts the choice of I'. The claim follows.

Following the claim, $\mathsf{OPT}_k(G) = \mathsf{OPT}_k(G')$. Moreover, $|V(G')| \leq k \cdot \Delta^{\leq 2}(G)$ by construction. Finally, it is straightforward to find G' in polynomial time by following the construction above. □

In the case of unit weights, a stronger lemma can be proven. In the shrinking model, it will imply later that either we can directly find an independent set of size k, or that due to the shrinking the number of input squares can be assumed to be at most $16k/\delta^2$. This is because the shrinking will allow us to argue that $\Delta^{\leq 1}(G) \leq 16/\delta^2$, for some appropriately chosen graph G.

Lemma 10. *Given an instance (G, w, k) of parameterized* MWIS *with $w(v) = 1$ for all $v \in V(G)$ and $|V(G)| > k \cdot \Delta^{\leq 1}(G)$, one can in polynomial time find an independent set in G of size $k = \mathsf{OPT}_k(G)$.*

Proof. Again, we iteratively define a sequence of graphs $G = G_0, G_1, \ldots, G_k$ and a sequence of vertices $v_1, \ldots, v_k \in V(G)$ as follows. Consider $i = 1, \ldots, k$ in order. To define G_i based on G_{i-1}, let v_i be any vertex in G_{i-1}, and let $G_i = G_{i-1} - N^{\leq 1}_{G_{i-1}}[v_i]$. As we will show next, G_i never becomes empty, and thus this sequence is well-defined.

Note that each G_i is an induced subgraph of G, and thus $N^{\leq 1}_{G_{i-1}}[v_i]$ is always of size at most $\Delta^{\leq 1}(G)$. By trivial induction, we have $|V(G_i)| \geq |V(G)| - i \cdot \Delta^{\leq 1}(G)$, which is always positive due to $|V(G)| > k \cdot \Delta^{\leq 1}(G)$. Hence, the construction can be performed for all the k steps, yielding distinct vertices $v_1, \ldots, v_k$.

Now observe that $\{v_1, \ldots, v_k\}$ is an independent set in G, because if v_j and v_i were adjacent for some $j > i$, then $v_j \in N^{\leq 1}_{G_{i-1}}[v_i]$, which combined with the fact that $v_j \neq v_i$ contradicts the choice of v_j. Hence, $\{v_1, \ldots, v_k\}$ can be returned. It is straightforward to perform the construction in polynomial time. □

4.3 Kernels on unit squares

To streamline the description of the kernels, we need the following definitions. The *center* of a rectangle R_i is the point $c_i = (x_i^{(1)} + \frac{g_i}{2}, y_i^{(1)} + \frac{h_i}{2})$. For some $\lambda \geq 0$, we say that a family $\mathcal{R}$ of rectangles is λ-*distant* if $||c_i - c_j||_\infty \geq \lambda$ for any distinct $R_i, R_j \in \mathcal{R}$. Also, if $\mathcal{R}$ is a family of squares and $\mathcal{Q} \subseteq \mathcal{R}$ is a subfamily of $\mathcal{R}$, then we say that a mapping $\phi\colon \mathcal{R} \to \mathcal{Q}$ is λ-*close* if for every $R_i \in \mathcal{R}$ with $R_j = \phi(R_i)$, we have $g_j \leq g_i$, $w_j \geq w_i$, and $||c_i - c_j||_\infty \leq \lambda$. Given such a mapping and a subfamily $\mathcal{S}$ of $\mathcal{R}$, we use $(\phi(\mathcal{S}))^{-\delta}$ to denote the δ-shrinking of the image of ϕ on $\mathcal{S}$.

The next lemma will prove very useful.

Lemma 11. *Let $\delta \in (0,1)$ and let $\mathcal{R}$ be a family of squares, each of side length at least M. Suppose further that $\phi\colon \mathcal{R} \to \mathcal{Q}$ is a $(\delta M/2)$-close mapping, for some subfamily $\mathcal{Q} \subseteq \mathcal{R}$. Then for every $R_i \in \mathcal{R}$, we have $R_j^{-\delta} \subseteq R_i$ where $R_j = \phi(R_i)$.*

Proof. Observe that

$$R_i = \{p :\ ||p - c_i||_\infty < g_i/2\} \quad \text{and} \quad R_j^{-\delta} = \{p :\ ||p - c_j||_\infty < (1/2 - \delta/2) \cdot g_j\}.$$

By the $(\delta M/2)$-closeness of ϕ, we have $g_j \leq g_i$ and $||c_i - c_j|| \leq (\delta/2) \cdot M$. Hence, for every $p \in R_j^{-\delta}$, we have

$$\begin{aligned} ||p - c_i||_\infty &\leq ||p - c_j||_\infty + ||c_i - c_j||_\infty \\ &< (1/2 - \delta/2) \cdot g_j + (\delta/2) \cdot M \leq (1/2 - \delta/2) \cdot g_i + (\delta/2) \cdot g_i = g_i/2. \end{aligned}$$

Since p was chosen arbitrarily, we conclude that $R_j^{-\delta} \subseteq R_i$. □

The following corollary is now almost immediate.

Corollary 12. *Let $\delta \in (0,1)$, let $\mathcal{R}$ is a family of squares, each of side length at least M, let $\mathcal{Q} \subseteq \mathcal{R}$ be a subfamily of $\mathcal{R}$, let $\phi\colon \mathcal{R} \to \mathcal{Q}$ be a $(\delta M/2)$-close mapping, and let $\mathcal{S} \subseteq \mathcal{R}$ be an independent subfamily of squares. Then ϕ is injective on $\mathcal{S}$, and $(\phi(\mathcal{S}))^{-\delta}$ is an independent set in $\mathcal{Q}^{-\delta}$ of weight not smaller than $w(\mathcal{S})$. Consequently, $\mathsf{OPT}_k(\mathcal{Q}^{-\delta}) \geq \mathsf{OPT}_k(\mathcal{R})$ for each non-negative integer k.*

Proof. Observe that if $R_j = \phi(R_i) = \phi(R_{i'})$ for some $R_j, R_i, R_{i'} \in \mathcal{R}$, then by Lemma 11 we have $R_j^{-\delta} \subseteq R_i \cap R_{i'}$, so in particular R_i and $R_{i'}$ overlap. Hence, ϕ must be injective on any independent subfamily $\mathcal{S}$. Moreover, Lemma 11 ensures that $(\phi(R_i))^{-\delta} \subseteq R_i$ for each $R_i \in \mathcal{S}$, so independence of $\mathcal{S}$ implies the independence of $(\phi(\mathcal{S}))^{-\delta}$. By the definition of a close mapping, each square is

mapped to a square of at least the same weight, and hence $w(\phi(\mathcal{S})) \geq w(\mathcal{S})$. The last assertion, that $\mathsf{OPT}_k(\mathcal{Q}^{-\delta}) \geq \mathsf{OPT}_k(\mathcal{R})$, follows immediately from the previous ones and the definition of $\mathsf{OPT}_k(\cdot)$. □

We now give our main tool for the construction of kernels for unit squares; it can be seen as a sparsification of $\mathcal{R}$.

Lemma 13. *Let k be a non-negative integer, let $\delta \in (0,1)$, and let $\mathcal{R}$ be a family of unit squares of non-uniform weight. Then one can in polynomial time construct a $(\delta/2)$-distant subfamily $\mathcal{Q} \subseteq \mathcal{R}$ together with a $(\delta/2)$-close mapping $\phi\colon \mathcal{R} \to \mathcal{Q}$.*

Proof. To construct $\mathcal{Q}$ and ϕ, we iteratively perform the following procedure. Take a square $R_j \in \mathcal{R}$ of maximum weight. Add R_j to $\mathcal{Q}$. For any square $R_i \in \mathcal{R}$ for which $||c_i - c_j||_\infty < (\delta/2)$, remove R_i from $\mathcal{R}$ and set $\phi(R_i) = R_j$. Iterate this procedure until $\mathcal{R}$ becomes empty. It is clear from the construction that $\mathcal{Q}$ constructed in this manner is $(\delta/2)$-distant, and $\phi\colon \mathcal{R} \to \mathcal{Q}$ is $(\delta/2)$-close; this is by the choice of R_j and the fact that all side lengths are equal. Moreover, the construction can be performed in polynomial time. □

We are ready to give a kernel for unit squares of uniform weight. In a nutshell, we use Lemma 13 to find a $(\delta/2)$-distant subfamily $\mathcal{Q}$ of $\mathcal{R}$. Then we transfer an argument made by Alber and Fiala [4, Corollary 3.1] to kernelize λ-distant families of unit disks to the setting of unit squares, and apply it to $\mathcal{Q}$ to obtain the final kernel.

Theorem 3. *Let k be a non-negative integer, let $\delta \in (0,1)$, and let $\mathcal{R}$ be a family of unit squares of uniform weight. Then there is a kernel for $\mathcal{R}$ and δ of size at most $16k/\delta^2$.*

Proof. Apply the algorithm of Lemma 13 to $\mathcal{R}$, thereby constructing a $(\delta/2)$-distant subfamily $\mathcal{Q} \subseteq \mathcal{R}$ and a $(\delta/2)$-close mapping $\phi\colon \mathcal{R} \to \mathcal{Q}$. By Corollary 12, we have $\mathsf{OPT}_k(\mathcal{Q}^{-\delta}) \geq \mathsf{OPT}_k(\mathcal{R})$. Consider the intersection graph G of $\mathcal{Q}^{-\delta}$.

First, suppose that $|\mathcal{Q}^{-\delta}| = |V(G)| > k \cdot \Delta^{\leq 1}(G)$. Then, we can apply Lemma 10 to G, thus in polynomial time finding an independent set I in G of size $k = \mathsf{OPT}_k(G)$. Then, we return the subfamily $\mathcal{S}$ of squares in $\mathcal{R}$ that corresponds to I. Note that $k = w(\mathcal{S}) = \mathsf{OPT}_k(\mathcal{Q}^{-\delta}) \geq \mathsf{OPT}_k(\mathcal{R})$ and that $\mathcal{S}^{-\delta}$ is an independent set. Hence, $\mathcal{S}$ satisfies the conditions of a kernel for $\mathcal{R}$ and δ of size at most k.

Second, suppose that $|\mathcal{Q}^{-\delta}| \leq |V(G)| \leq k \cdot \Delta^{\leq 1}(G)$. If we can bound $\Delta^{\leq 1}(G)$ by $(4/\delta)^2$, then $\mathcal{Q}$ satisfies the conditions of a kernel for $\mathcal{R}$ and δ of size at most $16k/\delta^2$, because $\mathcal{Q}$ is a subfamily of $\mathcal{R}$ and $\mathsf{OPT}_k(\mathcal{Q}^{-\delta}) \geq \mathsf{OPT}_k(\mathcal{R})$.

In order to bound $\Delta^{\leq 1}(G)$ by $(4/\delta)^2$, we prove that any square $R_i^{-\delta} \in \mathcal{Q}^{-\delta}$ overlaps at most $(4/\delta)^2$ squares of $\mathcal{Q}^{-\delta}$ (including itself). This implies the same bound on $\Delta^{\leq 1}(G)$. Observe that any square $R_j^{-\delta} \in \mathcal{Q}^{-\delta}$ that overlaps $R_i^{-\delta}$ has its center inside a square of side length $2 - 2\delta$ centered on c_i. Since $\mathcal{Q}$ is $(\delta/2)$-distant, so is $\mathcal{Q}^{-\delta}$, and thus open squares with side length $(\delta/2)$ centered at the centers of squares of $\mathcal{Q}^{-\delta}$ are pairwise disjoint. The union of all these open squares is contained in a square of side length $2 - \delta$ centered on c_i. Then a simple area bound shows that at most $(4/\delta)^2$ squares of $\mathcal{Q}^{-\delta}$ overlap R_i. The theorem follows. □

Theorem 4. *Let k be a non-negative integer, let $\delta \in (0,1)$, and let $\mathcal{R}$ be a family of unit squares of non-uniform weight. Then there is a kernel for $\mathcal{R}$ and δ of size at most $64k/\delta^2$.*

Proof. Again, apply the algorithm of Lemma 13 to $\mathcal{R}$, thereby constructing a $(\delta/2)$-distant subfamily $\mathcal{Q}$ and a $(\delta/2)$-close mapping $\phi\colon \mathcal{R} \to \mathcal{Q}$; in particular $\mathsf{OPT}_k(\mathcal{Q}^{-\delta}) \geq \mathsf{OPT}_k(\mathcal{R})$ by Corollary 12. Let G be the intersection graph of $\mathcal{Q}^{-\delta}$. We proceed identically as in the proof of Theorem 3, except that we replace our use of Lemma 10 by Lemma 9. More precisely, by applying the algorithm of

Lemma 9 on G, we compute an induced subgraph G' such that $\mathsf{OPT}_k(G') = \mathsf{OPT}_k(G) = \mathsf{OPT}_k(\mathcal{Q}^{-\delta})$ and $|V(G')| \leq k \cdot \Delta^{\leq 2}(G)$. This induced subgraph corresponds to a family of squares that forms a kernel. Hence, it suffices to bound $\Delta^{\leq 2}(G)$ by $64/\delta^2$. We prove that for any square $R_i^{-\delta} \in \mathcal{Q}^{-\delta}$, there are at most $64/\delta^2$ squares $R_j^{-\delta} \in \mathcal{Q}^{-\delta}$ that either overlap R_i, or that overlap a square of $\mathcal{Q}^{-\delta}$ that overlaps $R_i^{-\delta}$. Indeed, such $R_j^{-\delta}$ must have its center inside a square of side length $4 - 4\delta$ centered on c_i. Since $\mathcal{Q}$ is $(\delta/2)$-distant, a similar area bound as in Theorem 3 shows that at most $64/\delta^2$ squares of $\mathcal{Q}^{-\delta}$ overlap $R_i^{-\delta}$. □

4.4 Kernel for squares of non-uniform size and uniform weight

We first require the following auxiliary lemma, which superficially resembles Lemma 13. There is, however, a key difference: Lemma 13 can be applied to unit squares of non-uniform weights, while the next result treats squares of uniform weight and non-uniform size. This is reflected by the goal of the greedy choice in the proof.

Lemma 14. *Let k be a non-negative integer, let $\delta \in (0,1)$, and let $\mathcal{R}$ be a family of squares of uniform weight, each of side length at least M for some positive real M. Then one can in polynomial time construct a $(\delta M/2)$-distant subfamily $\mathcal{Q} \subseteq \mathcal{R}$ together with a $(\delta M/2)$-close mapping $\phi\colon \mathcal{R} \to \mathcal{Q}$.*

Proof. The proof is similar to that of Lemma 13, except that we use the side length of a square in order to create a $(\delta M/2)$-distant subfamily, rather than the weight. To construct $\mathcal{Q}$ and ϕ, we iteratively perform the following procedure. Take a square $R_j \in \mathcal{R}$ of minimum side length. Add R_j to $\mathcal{Q}$. For any square $R_i \in \mathcal{R}$ for which $||c_i - c_j||_\infty < (\delta M/2)$, remove R_i from $\mathcal{R}$ and set $\phi(R_i) = R_j$. Iterate this procedure until $\mathcal{R}$ becomes empty. Again, it is clear from the construction that $\mathcal{Q}$ constructed in this manner is $(\delta M/2)$-distant and $\phi\colon \mathcal{R} \to \mathcal{Q}$ is $(\delta M/2)$-close; this is by the choice of R_j and the fact that all weight are equal. Moreover, the construction can be performed in polynomial time. □

This lemma by itself leads to the following generalization of Theorem 3, which yields a kernel for instances in which the sizes of the input squares are in a bounded range. In a nutshell, we use Lemma 14 and then transfer an argument made by Alber and Fiala [4, Corollary 3.1] to kernelize λ-distant families of disks of bounded radius to our setting, just like Theorem 3.

Theorem 5. *Let k be a non-negative integer, let $\delta \in (0,1)$, let M_1, M_2 be real numbers such that $1 \leq M_1 \leq M_2$, and let $\mathcal{R}$ be a family of squares of uniform weight, each of side length between M_1 and M_2. Then there is a kernel for $\mathcal{R}$ and δ of size at most $16k/\delta^2 \cdot (M_2/M_1)^2$.*

Proof. Apply Lemma 14 to $\mathcal{R}$ and $M = M_1$, thereby obtaining a $(\delta M_1/2)$-distant subfamily $\mathcal{Q}$ and a $(\delta M_1/2)$-close mapping $\phi\colon \mathcal{R} \to \mathcal{Q}$; again, $\mathsf{OPT}_k(\mathcal{Q}^{-\delta}) \geq \mathsf{OPT}_k(\mathcal{R})$ by Corollary 12. Let G be the intersection graph of $\mathcal{Q}^{-\delta}$. We proceed identically as in the proof of Theorem 3, except for modifying the argumentation leading to the upper bound on $\Delta^{\leq 1}(G)$. Namely, observe that any square $R_j^{-\delta} \in \mathcal{Q}^{-\delta}$ that overlaps $R_i^{-\delta}$ has its center inside a square of side length $2M_2 - 2\delta M_1$ centered on c_i. Since $\mathcal{Q}$ is $M_1 \cdot (\delta/2)$-distant, a similar area bound as in Theorem 3 shows that at most $(4M_2/(M_1\delta))^2$ squares of $\mathcal{Q}$ overlap R_i. □

In order to extend this result to squares of arbitrary size, we need the following auxiliary lemma, which encapsulates the core combinatorial argument. It essentially states that if the sizes of two squares differ by a lot, then either one is contained in the other, or their shrunk counterparts are disjoint.

Lemma 15. *Let $\delta \in (0,1)$ and let R_i, R_j be any two squares for which $g_i \leq g_j$. Then R_i is contained in R_j, or $R_i^{-\delta}$ and $R_j^{-\delta}$ do not overlap, or $\frac{g_j}{g_i} < 2/\delta$.*

Proof. Suppose that R_i is not contained in R_j, and $R_i^{-\delta}$ and $R_j^{-\delta}$ overlap. We are going to prove that $\frac{g_j}{g_i} < 2/\delta$. Let ∂R_i and ∂R_j denote the boundaries of R_i and R_j respectively. Observe that R_i and R_j must overlap, since otherwise $R_i^{-\delta}$ and $R_j^{-\delta}$ would be disjoint. As R_i has not larger side length than R_j, and R_i is not contained in R_j, the boundaries ∂R_i and ∂R_j must intersect. Let x be any point in $\partial R_i \cap \partial R_j$.

Observe now that if $R_i^{-\delta}$ was contained in $R_j^{-\delta}$, then R_i would be contained in R_j, because $g_i \leq g_j$; this is contradicts the assumptions. Since $R_i^{-\delta}$ and $R_j^{-\delta}$ overlap, we conclude that their boundaries $\partial R_i^{-\delta}$ and $\partial R_j^{-\delta}$ also intersect. Let y be any point in $\partial R_i^{-\delta} \cap \partial R_j^{-\delta}$.

We now examine $||x-y||_\infty$, the ℓ_∞-distance between x and y. On the one hand, we have $x \in \partial R_i$ and $y \in \partial R_i^{-\delta} \subseteq R_i$. Hence,

$$||x-y||_\infty \leq g_i - \frac{\delta}{2} \leq g_i, \tag{4}$$

by noting that the ℓ_∞-distance between any point of ∂R_i and any point of $\partial R_i^{-\delta}$ is at most $g_i - \frac{\delta}{2} \leq g_i$.

On the other hand, $x \in \partial R_j$ and $y \in \partial R_j^{-\delta}$. Observe that each point of $\partial R_j^{-\delta}$ is at ℓ_∞-distance at least $\delta/2 \cdot g_j$ from any point of ∂R_j. Hence,

$$\delta/2 \cdot g_j \leq ||x-y||_\infty. \tag{5}$$

By combining (4) and (5), we conclude that $\frac{g_j}{g_i} < 2/\delta$. □

We can now give the kernel for squares of arbitrary size and uniform weight.

Theorem 6. *Let k be a non-negative integer, let $\delta \in (0,1)$, and let $\mathcal{R}$ be a family of squares of uniform weight and non-uniform size. Then there is a kernel for $\mathcal{R}$ and δ of size $\mathcal{O}(k^2 \cdot \frac{\log(1/\delta)}{\delta^3})$.*

Proof. The kernelization algorithm consists of several stages. In the first stage, we reduce $\mathcal{R}$ following a simple rule. If $\mathcal{R}$ contains two squares R_i, R_j such that R_i is contained in R_j, then in any subfamily $\mathcal{S}$ of $\mathcal{R}$ that is independent we can replace R_j by R_i; this replacement will not change the total weight, because the weights are uniform. Hence, removing R_j from $\mathcal{R}$ yields a family with the same value of $\mathsf{OPT}_k(\cdot)$. We apply this rule exhaustively. The resulting family, which we also call $\mathcal{R}$ by abuse of notation, has the same value of $\mathsf{OPT}_k(\cdot)$ as the original one, and thus the rule is safe. That is, from now on we may assume that for any two squares in $\mathcal{R}$, neither one square is contained in the other.

In the second stage, we partition the squares into levels according to their side length. For a non-negative integer ℓ, let

$$\mathcal{R}_\ell = \{R_i \in \mathcal{R} \mid (1+\delta)^\ell \leq g_i < (1+\delta)^{\ell+1}\}$$

be the family of squares of *level* ℓ. We call a level ℓ *active* if $\mathcal{R}_\ell \neq \emptyset$. Let $\gamma = 1 + \log_{1+\delta}(2/\delta)$. Thus, γ is essentially the $\log_{1+\delta}$ of the constant of Lemma 15, and note that $\gamma = \mathcal{O}(\log_{1+\delta}(1/\delta)) = \mathcal{O}(1/\delta \cdot \log 1/\delta)$.

Claim 1. *If $\mathcal{R}$ defines at least $\gamma \cdot k$ active levels, then one can find in polynomial time a subfamily $\mathcal{P}$ of $\mathcal{R}$ of size k such that $\mathcal{P}^{-\delta}$ is independent.*

Proof. By assumption, there are active levels $\ell_1, \ldots, \ell_k$ such that $\ell_j \geq \ell_i + \gamma$ for all $i, j \in \{1, \ldots, k\}$, $i < j$. Then

$$(1+\delta)^{\ell_j} \geq (1+\delta)^{\ell_i + 1 + \log_{1+\delta}(2/\delta)} \geq (1+\delta)^{\ell_i+1} \cdot (2/\delta).$$

Then for any two squares R_i, R_j in levels ℓ_i and ℓ_j respectively, where $i < j$, we have $\frac{g_j}{g_i} \geq 2/\delta$; this is by the definition of the levels. By Lemma 15, we infer that $R_i^{-\delta}$ and $R_j^{-\delta}$ have to be disjoint, because R_i is not contained in R_j by the assumption on $\mathcal{R}$ after the first stage.

For any $i \in \{1, \ldots, k\}$, let R_i be an arbitrary square of level ℓ_i; this is properly defined because ℓ_i is an active level. The preceding argument shows that if $\mathcal{P} = \{R_i \colon i \in \{1, \ldots, k\}\}$, then $\mathcal{P}^{-\delta}$ is an independent set of size and weight k. It now suffices to observe that $\mathcal{P}$ can be constructed in polynomial time by determining the level of each square, and counting the number of active levels. Since only the first $\log_{1+\delta} \max_{R_i \in \mathcal{R}} \{x_i^{(2)}, y_i^{(2)}\}$ levels may be active, this is indeed polynomial in the size of the input. ⌟

In the third stage, we consider two cases, depending on the number of active levels defined by $\mathcal{R}$. Suppose first that $\mathcal{R}$ defines at least $\gamma \cdot k$ active levels. Apply Claim 1, thereby constructing a subfamily $\mathcal{P} \subseteq \mathcal{R}$. Then $\mathsf{OPT}_k(\mathcal{P}^{-\delta}) = k \geq \mathsf{OPT}_k(\mathcal{R})$ by the properties of $\mathcal{P}$, following Claim 1 and the fact that all squares have uniform weight. Hence, $\mathcal{P}$ satisfies the conditions for a kernel for $\mathcal{R}$ and δ of size k, and the algorithm can output $\mathcal{P}$.

Suppose then that $\mathcal{R}$ defines less than $\gamma \cdot k$ active levels. For each active level ℓ, we apply the kernel of Theorem 5 to $\mathcal{R}_\ell$ with parameters k, $M_1 = (1+\delta)^\ell$, and $M_2 = (1+\delta)^{\ell+1}$. By inspection of the proof of Theorem 5 (and Theorem 3), this will either yield an independent subfamily $\mathcal{S}_\ell$ of $\mathcal{R}_\ell$ of size k, or provide a $(\delta M_1/2)$-distant subfamily $\mathcal{Q}_\ell \subseteq \mathcal{R}_\ell$ of size at most $16k(1+\delta)^2/\delta^2$, together with a $(\delta M_1/2)$-close mapping $f_\ell \colon \mathcal{R}_\ell \to \mathcal{Q}_\ell$. If the first case happens for some layer ℓ, then $|\mathcal{S}_\ell| = w(\mathcal{S}_\ell) = k = \mathsf{OPT}_k(\mathcal{R})$. Hence, $\mathcal{S}_\ell$ can be output as a kernel for $\mathcal{R}$ and δ of size k. If the first case never happens, that is, the second case happens for all active levels ℓ, then define $\mathcal{Q}$ to be the union of $\mathcal{Q}_\ell$ through all active levels ℓ. Then

$$|\mathcal{Q}| \leq (\gamma \cdot k) \cdot 16k(1+\delta)^2/\delta^2 = \mathcal{O}\left(k^2 \cdot \frac{\log(1/\delta)}{\delta^3}\right).$$

Moreover, $\mathcal{Q}$ can be computed in polynomial time.

Now in order to argue that $\mathcal{Q}$ is a kernel for $\mathcal{R}$ and δ, it suffices to show that $\mathsf{OPT}_k(\mathcal{Q}^{-\delta}) \geq \mathsf{OPT}_k(\mathcal{R})$. Take an independent subfamily $\mathcal{S} \subseteq \mathcal{R}$ with $|\mathcal{S}| = w(\mathcal{S}) = \mathsf{OPT}_k(\mathcal{R})$. Let us define $\mathcal{S}'$ to be the union of $f_\ell(\mathcal{S} \cap \mathcal{R}_\ell)$ for all active levels ℓ. Clearly $\mathcal{S}' \subseteq \mathcal{Q}$. By Lemma 11, f_ℓ maps each square $R_i \in \mathcal{S} \cap \mathcal{R}_\ell$ to a square $R_j = f_\ell(R_i)$ such that $R_j^{-\delta} \subseteq R_i$. Consequently, since $\mathcal{S}$ is independent, we infer that $(\mathcal{S}')^{-\delta}$ is also independent and $|\mathcal{S}'| = |\mathcal{S}|$. This means that $\mathsf{OPT}_k(\mathcal{Q}^{-\delta}) \geq |\mathcal{S}'| = |\mathcal{S}| = \mathsf{OPT}_k(\mathcal{R})$. Hence, $\mathcal{Q}$ indeed can be output as a kernel for $\mathcal{R}$ and δ of size $\mathcal{O}\left(k^2 \cdot \frac{\log(1/\delta)}{\delta^3}\right)$. □

4.5 Bit encodings

The algorithms of Theorems 3, 4, and 6 can be used to reduce the search space to a subfamily of size bounded in terms of k and $1/\delta$, when solving the appropriate subcases MWISR with δ-shrinking. However, in the field of kernelization, one often measures the size of a kernel by the bit size of its encoding. Definition 2 forbids changing the weights or the coordinates of the rectangles in the kernel. Hence, if we follow the definition, then we cannot bound the bit size of the encoding as a function of k and $1/\delta$. However, after computing the kernel, we can still discuss whether it is possible to replace the weights and the coordinates with numbers with a smaller encoding, so that the relevant combinatorial properties of the instance are preserved. We now show that this essentially can be done using known results, at least for the kernels obtained in this work.

First, we address the weights. Using the recent technique of Etscheid et al. [15], we can assume that the weight function for a kernel $\mathcal{Q}$ can be encoded using $\mathcal{O}(|\mathcal{Q}|^4)$ bits. Indeed, they show in [15, Example 1] that for any family X, weight function $w \colon X \to \mathbb{Q}$, and $W \in \mathbb{Q}$, one can in polynomial time find a weight function $w' \colon X \to \mathbb{Q}$ and $W' \in \mathbb{Q}$ of total encoding length $\mathcal{O}(|X|^4)$ such that $w(Y) \geq W$ if and only if $w'(Y) \geq W'$ for any $Y \subseteq X$. Hence, if we transform our parameterized setting to a parameterized decision setting, where OPT_k must be at least W, then we can find a

suitable encoding of the weight function. Our kernels do not rely on the a-priori knowledge of W, and therefore, we did not describe our kernels in the decision setting and remained more general instead.

Second, we address the representation of the squares. In Müller et al. [24], it is shown that any intersection graph of axis-parallel squares with n vertices has a representation using $\mathcal{O}(n \log n)$ or $\mathcal{O}(n^2)$ bits, depending on whether the squares have uniform or non-uniform size. However, actually finding a representation is NP-hard [5, 31].

Fortunately, the result of Müller et al. [24, Theorems 3 and 4] implies that when a representation of an intersection graph of (unit) squares is already given, then a representation with $\mathcal{O}(n \log n)$ or $\mathcal{O}(n^2)$ bits respectively can be found in polynomial time. We briefly describe the intuition. In the case of unit squares (see [24, Theorem 3]), a representation can be decomposed into two unit interval graphs by projecting on the axes, and a new representation for both these graphs can be computed in polynomial time using $\mathcal{O}(n \log n)$ bits. In the case of arbitrary squares (see [24, Theorem 4]), a representation can be transformed into a linear program, and a new representation can be found by using Cramer's rule to solve the linear program. This also shows that the NP-hardness of finding a representation is not in giving the actual numbers, but in finding the right "ordering" of the squares.

The above results imply that if a kernel uses a subfamily of size $f(k)$ of the squares in the input, which is the case for our kernels, then a representation of the squares in the kernel using $\mathcal{O}(f(k) \log f(k))$ or $\mathcal{O}(f^2(k))$ bits can be found in polynomial time, because the original squares imply a representation of the subfamily. An important caveat, however, is that finding such a new representation might change what squares are disjoint after shrinking by δ. Hence, if $\mathcal{Q}$ is the kernel, then we should rather find a new representation of $\mathcal{Q}^{-\delta}$. This makes sense if we think of the original motivation of the problem, which is to reduce the problem to solving MWISR in $\mathcal{Q}^{-\delta}$ (see also the remark below Definition 2).

4.6 Applications

We now present applications of our kernelization results for the design of FPT algorithms and EPTASes for the corresponding variants of MWISR. Recall that our kernelization results only apply to squares, and therefore, the applications in this section are significantly less general than Theorem 2 and Theorem 1. However, the resulting algorithms are substantially faster.

Throughout this section, for brevity we use the $\mathcal{O}_\rho(\cdot)$-notation that, for a parameter ρ, hides factors that depend only on ρ.

Parameterized algorithms. We obtain faster parameterized algorithms for squares of uniform weight or uniform size. In particular, we will obtain so-called *subexponential* algorithms, where the exponent depends sublinearly on k. The design of such algorithms in the planar and geometric setting is an important direction in parameterized complexity research, see e.g. [12, 16, 20, 23, 25, 26]. We note that the exponent of Theorem 1, which applies to arbitrary rectangles, is polynomial in k. Hence, Theorem 7 and 8 below are a major improvement whenever all the input rectangles are squares of uniform weight or of uniform size.

We present two different and incomparable results. The first is obtained by combining our kernelization with the algorithm of Marx and Pilipczuk [23], and for the second we similarly use the results of Alber and Fiala [4].

Theorem 7. *There is an algorithm that given a weighted family $\mathcal{R}$ of n axis-parallel squares of total encoding size N and of uniform size or of uniform weight, and parameters k and δ, runs in time $(k/\delta)^{\mathcal{O}(\sqrt{k})} \cdot (nN)^c$ for some constant c, and outputs a subfamily $\mathcal{S} \subseteq \mathcal{R}$ such that $|\mathcal{S}| \leq k$, $\mathcal{S}^{-\delta}$ is independent, and $w(\mathcal{S}) \geq \mathsf{OPT}_k(\mathcal{R})$.*

Proof. We apply a result of Marx and Pilipczuk [23]. They gave an algorithm that, given a family $\mathcal{D}$ of n polygons of non-uniform weight[1] in the plane and a parameter k, finds a subfamily of polygons that is independent and has size at most k, and which achieves the maximum possible weight subject to these constraints. The running time of the algorithm is $|\mathcal{D}|^{\mathcal{O}(\sqrt{k})} \cdot (nN)^{\mathcal{O}(1)}$, where N is the total bit size of the input. By applying this algorithm on $\mathcal{Q}^{-\delta}$, for a kernel $\mathcal{Q}$ obtained by applying any of Theorem 3, 4, or 6, we can compute $\mathsf{OPT}_k(\mathcal{Q}^{-\delta})$ in time $|\mathcal{Q}^{-\delta}|^{\mathcal{O}(\sqrt{k})} \cdot (nN)^{\mathcal{O}(1)} = (k/\delta)^{\mathcal{O}(\sqrt{k})} \cdot (nN)^{\mathcal{O}(1)}$. □

Theorem 8. *There is an algorithm that given a weighted family $\mathcal{R}$ of n axis-parallel squares of total encoding size N, uniform size, and non-uniform weight, and parameters k and δ, runs in time $2^{\mathcal{O}_\delta(\sqrt{k})} \cdot (nN)^c$ for some constant c, and outputs a subfamily $\mathcal{S} \subseteq \mathcal{R}$ such that $|\mathcal{S}| \leq k$, $\mathcal{S}^{-\delta}$ is independent, and $w(\mathcal{S}) \geq \mathsf{OPT}_k(\mathcal{R})$.*

Proof. Let $\mathcal{Q}$ be the family returned by Theorem 4 applied to k, δ, and $\mathcal{R}$. Observe that $\mathcal{Q}$ can be computed in polynomial time and that $|\mathcal{Q}| = \mathcal{O}_\delta(k)$. Moreover, an inspection of the proof of Theorem 4 reveals that $\mathcal{Q}$ is $(\delta/2)$-distant, and thus the same holds for $\mathcal{Q}^{-\delta}$. A slight adaptation of a result of Alber and Fiala [4, Corollary 4.1] (see also [29, Theorem 6.2.4]) gives an algorithm that, given a family $\mathcal{D}$ of n λ-distant unit squares of non-uniform weight[2] in the plane, finds a subfamily of $\mathcal{D}$ that is independent and achieves the maximum possible weight subject to this constraint. The running time of the algorithm is $2^{\mathcal{O}_\lambda(\sqrt{|\mathcal{D}|})} \cdot (nN)^{\mathcal{O}(1)}$, where N is the total bit size of the input. By applying this algorithm on $\mathcal{Q}^{-\delta}$, we can compute $\mathsf{OPT}_k(\mathcal{Q}^{-\delta})$ in time $2^{\mathcal{O}_\delta(\sqrt{k})} \cdot (nN)^{\mathcal{O}(1)}$. □

We can also obtain a faster algorithm for squares of uniform weight and non-uniform size by combining the results of Alber and Fiala [4] with the kernel of Theorem 6 and a later lemma (Lemma 17) about the ply of the objects. However, the running time obtained in this way is $2^{\mathcal{O}_\delta(k)} \cdot (nN)^c$, which is worse than what we can obtain through Theorem 7.

Approximation algorithms. Next, we show how to obtain faster approximation algorithms for squares of uniform weight or of uniform size. In order to prove these results, we need several auxiliary steps, and the following notion in particular. The *ply* of a family of geometric objects in the two-dimensional plane is defined as the maximum number of objects in the family that contain any point in the plane. This is sometimes also called the *maximum depth* of any point. This enables us to state following auxiliary lemma, which transfers a result of Alber and Fiala [4, Lemma 4.1] from disks to squares.

Lemma 16. *Let $\lambda > 0$, let M be a real number, and let $\mathcal{R}$ be a λ-distant family of squares, each of side length at most M. Then $\mathcal{R}$ has ply $\mathcal{O}(M^2/\lambda^2)$.*

Proof. This follows using essentially the same area bound as in Theorem 5 and 3; see also the work of Alber and Fiala [4, Lemma 4.1]. □

Now we can prove the following lemma.

Lemma 17. *Let k be a non-negative integer, let $\delta \in (0, 1)$, let $\mathcal{R}$ be a family of squares of uniform weight and non-uniform size, and let $\mathcal{Q} \subseteq \mathcal{R}$ be the family returned by Theorem 6 when applied to k, δ, and $\mathcal{R}$. Then the ply of $\mathcal{Q}^{-\delta}$ is $\mathcal{O}_\delta(1)$.*

[1] Marx and Pilipczuk state this result explicitly for the unweighted setting (Theorem 1.5 in [23]), but the result for the weighted setting follows from the most general statement (Theorem 3.1 in [23]).

[2] Alber and Fiala state this result explicitly for the unweighted setting [4, Corollary 4.1], but the result for the weighted setting follows in exactly the same way. Moreover, instead of their Theorem 4.2, we need to rely on a separator theorem by Smith and Wormald [27, Theorem 26]. See also [29, Theorem 6.2.4].

Proof. If $\mathcal{Q}^{-\delta}$ is an independent set, then $\mathcal{Q}^{-\delta}$ has ply at most 1. Hence, we can focus on the final case of the proof of Theorem 6, i.e., when the kernel of Theorem 5 is applied to each active level separately, for this is the only case when the algorithm does not return an independent set of size k. We adopt the notation from the proof of Theorem 5.

Let p be any point in the plane. We claim that any two squares of $\mathcal{Q}^{-\delta}$ that contain p differ by $\mathcal{O}(\log_{1+\delta}(2/\delta))$ levels (recall that levels were defined w.r.t. the original squares from $\mathcal{Q}$, but we consider their corresponding shrunken squares in $\mathcal{Q}^{-\delta}$). Let $R_i^{-\delta} \in \mathcal{Q}^{-\delta}$ be any square that contains p. We now consider any other square $R_j^{-\delta} \in \mathcal{Q}^{-\delta}$ that contains p. Assume w.l.o.g. that $g_i \leq g_j$ (the other case is symmetric). By the construction of $\mathcal{Q}$, R_i cannot contain R_j or vice versa. Since $R_i^{-\delta}$ and $R_j^{-\delta}$ overlap (in p), it follows from Lemma 15 that $\frac{g_j}{g_i} < 2/\delta$. Since levels are delimited by powers of $(1+\delta)$, the claim follows.

Now observe that each level of $\mathcal{Q}^{-\delta}$ has ply $\mathcal{O}_\delta(1)$ by Lemma 16, because by construction, each level ℓ of $\mathcal{Q}^{-\delta}$ is a family of $(\delta(1+\delta)^\ell/2)$-distant squares, each of side length at most $(1+\delta)^{\ell+1}$. Combined with the claim, this implies that the ply of $\mathcal{Q}^{-\delta}$ is $\mathcal{O}_\delta(1)$. □

The next theorem gives our EPTAS for axis-parallel squares. The existence of an EPTAS already follows from Theorem 2, but we state it because the obtained dependence of the running time on ϵ is far better than that of Theorem 2.

Theorem 9. *There is an algorithm that given a weighted family $\mathcal{R}$ of n axis-parallel squares with total encoding size N, of uniform size or of uniform weight, and parameters δ, ϵ, runs in time $2^{\mathcal{O}_\delta(1/\epsilon)} \cdot (nN)^c$ for some constant c, and outputs a subfamily $\mathcal{S} \subseteq \mathcal{R}$ such that $\mathcal{S}^{-\delta}$ is independent, and $w(\mathcal{S}) \geq (1-\epsilon)\mathsf{OPT}(\mathcal{R})$.*

Proof. For each $k = 1, \ldots, n$, we will construct a family $\mathcal{Q}_k \subseteq \mathcal{R}_k$ and a family $\mathcal{S}_k \subseteq \mathcal{Q}_k$ such that $\mathsf{OPT}_k(\mathcal{Q}_k^{-\delta}) \geq \mathsf{OPT}_k(\mathcal{R})$, $\mathcal{S}_k^{-\delta}$ is independent, and $w(\mathcal{S}_k) \geq (1-\epsilon)\mathsf{OPT}(\mathcal{Q}_k^{-\delta})$. Then

$$w(\mathcal{S}_k) \geq (1-\epsilon)\mathsf{OPT}(\mathcal{Q}_k^{-\delta}) \geq (1-\epsilon)\mathsf{OPT}_k(\mathcal{Q}_k^{-\delta}) \geq (1-\epsilon)\mathsf{OPT}_k(\mathcal{R}).$$

Since there is a value of k for which $\mathsf{OPT}_k(\mathcal{R}) = \mathsf{OPT}(\mathcal{R})$, the heaviest family $\mathcal{S}_k$ over all k satisfies the conditions of the theorem, because each $\mathcal{S}_k$ is an independent set. Hence, it suffices to find the families $\mathcal{Q}_k$ and $\mathcal{S}_k$ for each $k = 1, \ldots, n$.

The construction of $\mathcal{Q}_k$ and $\mathcal{S}_k$ depends on whether the squares in $\mathcal{R}$ have uniform size or uniform weight. Consider the case when the squares have uniform size. First, apply the kernel of Theorem 4 to k, δ, and $\mathcal{R}$, and let $\mathcal{Q}_k \subseteq \mathcal{R}$ denote the resulting family. By inspection of the proof of Theorem 4, $\mathcal{Q}_k$ is $(\delta/2)$-distant. Now recall that a result of Hunt et al. [19, Theorem 5.2] shows that Maximum Weight Independent Set of Rectangles has a $(1-\epsilon)$-approximation in time $2^{\mathcal{O}_\lambda(1/\epsilon)} \cdot N^{\mathcal{O}(1)}$ whenever all the input rectangles are λ-distant squares of uniform size and non-uniform weight[3]. We apply this result to the family $\mathcal{Q}_k^{-\delta}$, which is also $(\delta/2)$-distant, and ϵ. Let $\mathcal{S}_k \subseteq \mathcal{Q}_k$ be the subfamily of $\mathcal{Q}_k$ such that $\mathcal{S}_k^{-\delta}$ is the output of the approximation algorithm. Note that $\mathcal{Q}_k$ and $\mathcal{S}_k$ are as required, because $\mathsf{OPT}_k(\mathcal{Q}_k^{-\delta}) \geq \mathsf{OPT}_k(\mathcal{R})$ by Theorem 4, $\mathcal{S}_k^{-\delta}$ is independent by construction, and $w(\mathcal{S}_k) \geq (1-\epsilon)\mathsf{OPT}(\mathcal{Q}_k^{-\delta})$ by construction.

Consider now the case when the input squares have uniform weight. First, apply the kernel of Theorem 6 to k, δ, and $\mathcal{R}$, and let $\mathcal{Q}_k \subseteq \mathcal{R}$ denote the resulting family. By Lemma 17, $\mathcal{Q}_k^{-\delta}$ has ply $\mathcal{O}_\delta(1)$. Now recall that a result of Chan [8, Theorem 3.1] (see also [29, Section 7.3/7.4] for a generalization) shows that Maximum Weight Independent Set of Rectangles has a $(1-\epsilon)$-approximation in time $2^{\mathcal{O}_\tau(1/\epsilon)} \cdot N^{\mathcal{O}(1)}$ whenever all the input rectangles are squares of

[3]Hunt et al. state this result explicitly for the unweighted setting [19, Theorem 5.2], but the result for the weighted setting follows in exactly the same way. See also [29, Section 6.3.1/6.3.5] for a generalization and a better dependence on λ.

uniform weight, non-uniform size, and ply τ. We apply this result to the family $\mathcal{Q}_k^{-\delta}$. Let $\mathcal{S}_k \subseteq \mathcal{Q}_k$ be the subfamily of $\mathcal{Q}_k$ such that $\mathcal{S}_k^{-\delta}$ is the output of the approximation algorithm. Note that $\mathcal{Q}_k$ and $\mathcal{S}_k$ are as required, because $\mathsf{OPT}_k(\mathcal{Q}_k^{-\delta}) \geq \mathsf{OPT}_k(\mathcal{R})$ by Theorem 6, $\mathcal{S}_k^{-\delta}$ is independent by construction, and $w(\mathcal{S}_k) \geq (1-\epsilon)\mathsf{OPT}(\mathcal{Q}_k^{-\delta})$ by construction. □

5 Conclusions

In this paper, we have initiated the study of the shrinking model for parameterized geometric Independent Set, by giving FPT algorithms and polynomial kernels for the most basic variants. Most importantly, we have showcased that the shrinking model leads to robust tractability of problems that without this relaxation are hard from the parameterized perspective. We hope that this is the start of an interesting direction in the research on combinatorial optimization for geometric problems, as our work raises several concrete open problems. Can our FPT algorithm and EPTAS for axis-parallel rectangles (Theorems 2 and 1) be generalized to arbitrary convex polygons, as is the case for the PTAS [30]? Is it possible to improve the running time of said FPT algorithm to subexponential, that is, $2^{o(k)} \cdot (nN)^{\mathcal{O}(1)}$ for every fixed δ? What about polynomial kernels, that is, kernelization procedures with polynomial output guarantees, for more complex objects than squares? Here, it seems that our arguments apply mutatis mutandis to e.g. (unit) disks instead of (unit) squares, but it is conceivable that even the general setting of convex or star-convex polygons can be treated, for an appropriate definition of shrinking. Also, is there a polynomial kernel for squares of non-uniform size and non-uniform weight? This is not addressed in its full generality by our kernelization algorithms. Finally, can one give limits to tractability in the shrinking model, by showing W[1]-hardness or the nonexistence of polynomial kernels? We hope that the future work will address these, as well as multiple other open problems arising from this work.

Acknowledgements. The first author thanks Dániel Marx for helpful discussions about the setting.

References

[1] A. Adamaszek, P. Chalermsook, and A. Wiese. How to Tame Rectangles: Solving Independent Set and Coloring of Rectangles via Shrinking. In *APPROX/RANDOM 2015*, volume 40 of *LIPIcs*, pages 43–60. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2015.

[2] A. Adamaszek and A. Wiese. Approximation schemes for maximum weight independent set of rectangles. In *FOCS 2013*, pages 400–409. IEEE, 2013.

[3] A. Adamaszek and A. Wiese. A QPTAS for maximum weight independent set of polygons with polylogarithmically many vertices. In *SODA 2014*, pages 645–656. SIAM, 2014.

[4] J. Alber and J. Fiala. Geometric separation and exact solutions for the parameterized independent set problem on disk graphs. *J. Algorithms*, 52(2):134–151, 2004.

[5] H. Breu. *Algorithmic Aspects of Constrained Unit Disk Graphs*. PhD thesis, University of British Columbia, 1996.

[6] P. Chalermsook and J. Chuzhoy. Maximum independent set of rectangles. In *SODA 2009*, pages 892–901. SIAM, 2009.

[7] T. M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. *Journal of Algorithms*, 46(2):178–189, 2003.

[8] T. M. Chan. A note on maximum independent sets in rectangle intersection graphs. *Information Processing Letters*, 89(1):19–23, 2004.

[9] T. M. Chan and S. Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. *Discrete & Comp. Geometry*, 48(2):373–392, 2012.

[10] J. Chuzhoy and A. Ene. On approximating maximum independent set of rectangles. In *FOCS 2016*, pages 820–829. IEEE, 2016.

[11] M. Cygan, F. V. Fomin, L. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, 2015.

[12] E. D. Demaine, F. V. Fomin, M. T. Hajiaghayi, and D. M. Thilikos. Subexponential parameterized algorithms on bounded-genus graphs and h-minor-free graphs. *J. ACM*, 52(6):866–893, 2005.

[13] R. G. Downey and M. R. Fellows. Fixed-parameter tractability and completeness II: On completeness for W[1]. *Theoretical Computer Science*, 141(1):109–131, 1995.

[14] T. Erlebach, K. Jansen, and E. Seidel. Polynomial-time approximation schemes for geometric intersection graphs. *SIAM J. on Computing*, 34(6):1302–1323, 2005.

[15] M. Etscheid, S. Kratsch, M. Mnich, and H. Röglin. Polynomial kernels for weighted problems. In *MFCS 2015*, volume 9235 of *LNCS*, pages 287–298. Springer, 2015.

[16] F. V. Fomin, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. Subexponential parameterized algorithms for planar and apex-minor-free graphs via low treewidth pattern covering. In *FOCS 2016*, pages 515–524. IEEE, 2016.

[17] S. Har-Peled. Quasi-polynomial time approximation scheme for sparse subsets of polygons. In *SOCG 2014*, pages 120–129, 2014.

[18] D. S. Hochbaum and W. Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *J. ACM*, 32(1):130–136, 1985.

[19] H. B. Hunt III, M. V. Marathe, V. Radhakrishnan, S. S. Ravi, D. J. Rosenkrantz, and R. E. Stearns. NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. *J. Algorithms*, 26(2):238–274, 1998.

[20] P. N. Klein and D. Marx. A subexponential parameterized algorithm for subset TSP on planar graphs. In *SODA*, pages 1812–1830. SIAM, 2014.

[21] A. Lingas and M. Wahlen. A note on maximum independent set and related problems on box graphs. *Inf. Process. Lett.*, 93(4):169–171, 2005.

[22] D. Marx. On the optimality of planar and geometric approximation schemes. In *FOCS 2007*, pages 338–348. IEEE Computer Society, 2007.

[23] D. Marx and M. Pilipczuk. Optimal parameterized algorithms for planar facility location problems using Voronoi diagrams. In *ESA 2015*, volume 9294 of *LNCS*, pages 865–877. Springer, 2015.

[24] T. Müller, E. J. van Leeuwen, and J. van Leeuwen. Integer representations of convex polygon intersection graphs. *SIAM J. Discrete Math.*, 27(1):205–231, 2013.

[25] M. Pilipczuk, M. Pilipczuk, P. Sankowski, and E. J. van Leeuwen. Subexponential-time parameterized algorithm for Steiner Tree on planar graphs. In *STACS 2013*, volume 20 of *LIPIcs*, pages 353–364. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2013.

[26] M. Pilipczuk, M. Pilipczuk, P. Sankowski, and E. J. van Leeuwen. Network sparsification for Steiner problems on planar and bounded-genus graphs. In *FOCS 2014*, pages 276–285. IEEE Computer Society, 2014.

[27] W. D. Smith and N. C. Wormald. Geometric separator theorems & applications. In *FOCS 1998*, pages 232–243. IEEE Computer Society, 1998.

[28] E. J. van Leeuwen. Better approximation schemes for disk graphs. In *SWAT 2006*, volume 4059 of *Lecture Notes in Computer Science*, pages 316–327. Springer, 2006.

[29] E. J. van Leeuwen. *Optimization and Approximation on Systems of Geometric Objects*. PhD thesis, University of Amsterdam, 2009.

[30] A. Wiese. Independent set of convex polygons: From n^ϵ to $1+\epsilon$ via shrinking. In *LATIN 2016*, pages 700–711. Springer, 2016.

[31] M. Yannakakis. The complexity of the partial order dimension problem. *SIAM J. Algebraic Discr. Methods*, 3(3):351–358, 1982.

[32] D. Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3:103–128, 2007.